Int J Comput Vis (2013) 103:22–59
DOI 10.1007/s11263-012-0592-x

Toward a Comprehensive Framework for the Spatiotemporal Statistical Analysis of Longitudinal Shape Data

Stanley Durrleman · Xavier Pennec · Alain Trouvé · José Braga · Guido Gerig · Nicholas Ayache

Received: 12 August 2011 / Accepted: 17 October 2012 / Published online: 13 November 2012
© Springer Science+Business Media New York 2012

Abstract This paper proposes an original approach for the statistical analysis of longitudinal shape data. The proposed method allows the characterization of typical growth patterns and subject-specific shape changes in repeated time-series observations of several subjects. This can be seen as the extension of usual longitudinal statistics of scalar measurements to high-dimensional shape or image data. The method is based on the estimation of continuous subject-specific growth trajectories and the comparison of such temporal shape changes across subjects. Differences between growth trajectories are decomposed into morphological deformations, which account for shape changes independent of the time, and time warps, which account for different rates of shape changes over time. Given a longitudinal shape data set, we estimate a mean growth scenario representative of the population, and the variations of this scenario both in terms of shape changes and in terms of change in growth speed. Then, intrinsic statistics are derived in the space of spatiotemporal deformations, which characterize the typical variations in shape and in growth speed within the studied population. They can be used to detect systematic developmental delays across subjects. In the context of neuroscience, we apply this method to analyze the differences in the growth of the hippocampus in children diagnosed with autism, developmental delays and in controls. Result suggest that group differences may be better characterized by a different speed of maturation rather than shape differences at a given age. In the context of anthropology, we assess the differences in the typical growth of the endocranium between chimpanzees and bonobos. We take advantage of this study to show the robustness of the method with respect to change of parameters and perturbation of the age estimates.

Keywords Longitudinal data · Statistics · Growth · Shape regression · Spatiotemporal registration · Time warp

Electronic supplementary material The online version of this article (doi:10.1007/s11263-012-0592-x) contains supplementary material, which is available to authorized users.

S. Durrleman (✉)· G. Gerig
Scientific Computing and Imaging (SCI) Institute,
72 S. Central Drive, Salt Lake City, UT 84112, USA
e-mail: stanley.durrleman@gmail.com, stanley@sci.utah.edu

S. Durrleman · X. Pennec · N. Ayache
Asclepios team-project, INRIA Sophia Antipolis, 2004 route des Lucioles, 06902 Sophia Antipolis, France

S. Durrleman · A. Trouvé
Centre de Mathématiques et Leurs Applications (CMLA), CNRS-ENS Cachan, 61 avenue du Président Wilson, 94235 Cachan, France

J. Braga
Laboratoire de paléoanthropologie assistée par ordinateur, CNRS-Université de Toulouse (Paul Sabatier), 37 allées Jules Guesde, 37073 Toulouse, France

1 Spatiotemporal Variability of Longitudinal Data

Many scientific questions can be expressed in terms of changes or alterations of a dynamical process. In camera surveillance, one aims at distinguishing normal from abnormal behaviors in video sequences. In clinical studies, one wants to characterize anatomical or functional changes due to disease progression, clinical intervention or therapy. In neuroscience, one studies the neurodevelopment or the neurodegeneration of the brain and its related structures. In cardiac imaging, one looks for abnormal patterns in the heart motion. What make these questions so challenging is that the evolving object of interest changes in appearance in different situations. In video

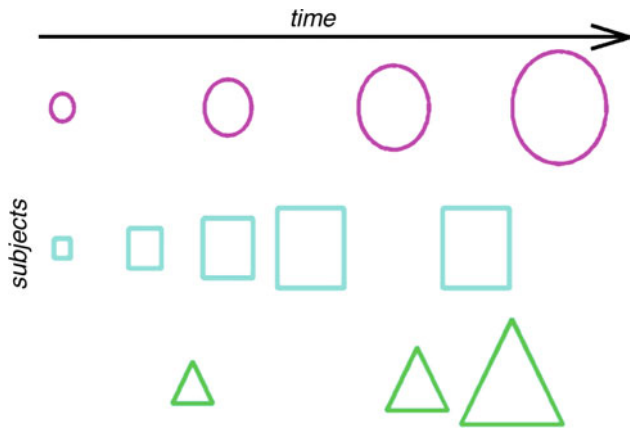


Fig. 1 Synthetic example of a longitudinal data set with 3 subjects. Each subject has been observed a few times and at different time-points. The aim of the spatiotemporal variability analysis is to describe the variability of this population in two ways: the geometrical variability (there is a *circle*, a *square* and a *triangle*), and the variability in terms of change of dynamics of evolution (for instance, the *square* grows first at a faster rate than the *circle* and then slows down.)

sequences for instance, we want to distinguish a normal from an abnormal behavior behind the large variety of the shapes and the motions of the silhouettes. Similarly, every brain has a different shape, whereas its maturation may follow some common patterns that we would like precisely to describe and quantify.

From the point of view of data analysis and pattern theory, these problems can be addressed by the statistical analysis of longitudinal data sets. A longitudinal data set consists of the observation of a set of homologous objects (such as silhouettes of people or anatomical structures), each object being observed repeatedly at several time points. An abstract example of such a data set is given in Fig. 1, which illustrates the sampling of individual growth trajectories of different subjects. The analysis of such longitudinal data sets should lead to the qualitative and quantitative assessment of change trajectories, to the detection of common growth patterns shared in a population, and to the characterization of their appearances in different subjects.

Longitudinal analysis differs from the usual cross-sectional variability analysis in that it takes into account the inherent correlation of repeated measurements of the same individuals. It must also provide a model of how an individual subject's trajectory changes relative to another subject. At the population level, we typically analyze how the subjects are distributed within a group by estimating a mean configuration and its variance. For longitudinal data, the mean configuration may be a "mean growth scenario", which averages the growth patterns in the population. The analysis of its variance explains how each subject's trajectory differs from the mean growth scenario. Such a statistical approach based on mean and variance is well-known for scalar measurements and for analysis of cross-sectional shape data, for which the mean is usually called "template" or "atlas". The extension of these concepts for longitudinal shape data is challenging, as no consensus has emerged about how to combine shape changes over time and shape changes across subjects.

In this paper, we propose a consistent conceptual and computational framework to address these questions: (i) the estimation of subject-specific trajectories via the introduction of a growth model as a smooth deformation of the baseline shape, (ii) the comparison of different trajectories via spatiotemporal mappings which align both the shape of different subjects *and* the tempo of their respective evolution, (iii) the estimation of a "mean growth scenario" representative of a given population, and (iv) the statistical analysis of the typical variations of this mean scenario in the studied population. The proposed methodology does not require that the subjects are observed with the same number of samples or at the same time-points.

One of the main contributions of this methodology is that it models the changes in individual trajectories both as morphological changes, which account for the different appearances of the object, and as dynamical changes, which account for different paces of evolution. At the population level, this assumes that the development of different subjects shares the same growth patterns, up to changes in shape and changes in the tempo of the development. This enables in particular the characterization of the effect of a pathology as a systematic developmental delays in the growth of a given organ.

The detailed explanation of the method and its related algorithms is given in Sect. 3. Section 2 explains how the proposed framework consistently embeds different concepts introduced in the literature and highlights different possible modeling choices. In Sect. 4, we show how the method can be used to characterize the effect of autism and developmental delay in the growth speed of the hippocampus. In Sect. 5, the method will be used to quantitatively assess the relative developmental delay of the endocranial growth between bonobos and chimpanzees. We will show that this estimation is robust to parameter changes and changes in the age estimates of the samples.

2 An Emerging Framework for the Analysis of Longitudinal Shape Data

This section presents a survey of the literature on the topic of longitudinal analysis of shape data. This will highlight which tools and concepts need to be linked into a common statistical framework. We will also make clear that different modeling choices are possible. We will propose to follow the approach that seems the more adapted to the targeted applications.

2.1 Previous Research to Design 4D Statistical Analysis

2.1.1 4D Analysis Meant as Regression or Tracking

The first kind of so-called 4D-analysis proposes to estimate a continuous sequence from a set of time-indexed shapes or images of the same subject. In Mansi et al. (2009), one estimates a cross-sectional atlas from time series data and then analyzes the correlations between the modes of variability and the age of the subject, considered as an explanatory variable. These correlations may be used to estimate a synthetic growth scenario for a given individual. Other approaches, which do not rely on a cross-sectional atlas, include work by de Craene et al. (2009), in which the authors use Large Diffeomorphic Free Form Deformations to estimate time-varying deformations between the first and the last sample of a sequence of images. In the same spirit, Davis et al. (2007) propose to perform the regression of a sequence of images via a generalization of the kernel regression method to Riemannian manifolds. Growth scenarios could also be estimated based on stochastic growth models as in Grenander et al. (2007) and Trouvé and Vialard (2010) or on twice differentiable flows of deformations as in Fishbaugh et al. (2011).

These methods are pure regression methods. If they are used with several subjects scanned several times, these regression methods return a single evolution, the most probable evolution in some sense. They do not take into account that data at different time points may come from the same subject or from different subjects. It averages shape evolutions without discarding the inter-subject variability, which leads to "fuzzy" estimation like the average of a set of non-registered images. By contrast, in Thompson et al. (2000) and Gogtay et al. (2008), registrations between baseline and follow-up scans of the same subject are performed and the evolutions of *scalar* measurements extracted from the registration are compared across subjects. A main contribution of our paper will be to extend this framework for scalar measurements to the high-dimensional space of shapes. In Khan and Beg (2008), the authors propose to perform a regression of the image sequence of every subject separately and then to average the time dependent velocity field of each regression to estimate a typical scenario of evolution. This approach is limited to situations where each sequence is registered in the same reference frame, but no details of how to perform registration of time-indexed sequences of images is given.

2.1.2 3D-Registration of 4D-Sequences

The problem of registering individual trajectories has been investigated in different communities. In Chandrashekara et al. (2003), the motion of the heart of each subject is tracked through time. Then the registration between the baseline image of two subjects is used to transport the velocity field of the tracking from one subject's space to the other. This approach could also include the estimation of a template image at the baseline time-point using usual cross-sectional atlas construction methods, like in Ehrhardt et al. (2008) and Qiu et al. (2008, 2009). All these methods assume that the inter-subject variability can be captured considering only the baseline images. Using these deformations for registering the whole time-indexed sequence of images is arguable, since they do not take into account anatomical features which may appear later in the sequence.

This issue has been addressed in Peyrat et al. (2008) who proposed to register a time-indexed sequence by computing deformations between any pair of successive scans of the same subject and between any pair of scans of two different subjects at the same time-point. Such an approach takes all temporal information into account and therefore leads to a much more robust registration scheme. However, this method assumes that every time-indexed sequence has exactly the same number of images, which are acquired at time points which correspond across subjects. By contrast, in longitudinal studies, only a few scans per subject are available, and the number of scans may vary for different subjects. This issue has been addressed in Hart et al. (2010), who proposed an interpolation scheme to average individual trajectories at every time-point independently. However, none of these approaches take into account the inherent temporal correlations between successive inter-subject registration. From a statistical point of view, this means that the inter-subject variability at two different time points are considered as independent variables. As the sampling of the image sequence becomes finer and finer, the number of variables to estimate becomes larger and larger. A main contribution of this paper will be precisely to define a generative statistical model, which takes into account the temporal correlations between inter-subject registrations at different time points, and to provide a way to estimate these correlations from a finite set of observations.

In Gerig et al. (2006), a template image is built at every time-point independently. Then, the deformations between the baseline scan and the follow-up of any subjects are compared to the deformation between the baseline atlas and the follow-up atlas. This approach focuses on the analysis of the cross-sectional variability over time. However, since the template image is built at each time point independently, it is not clear whether the difference between the baseline atlas and the follow-up atlas is the average of every subject's evolution. Moreover, such a method requires that the distribution of age in the longitudinal data set is clustered at two distinct ages, which is a special case. A similar approach has also been proposed in Aljabar et al. (2008).

2.1.3 Taking into Account Temporal Re-alignment

The methods cited previously propose a way to combine the subject-specific growth with the inter-subject variability: time-series image sequence are processed by a combination of 3D deformations. In particular, the age at which the subjects are scanned is considered an absolute time which corresponds across subjects. This assumes that at a given age, every subject is at the same development stage and that their anatomy can be compared. Such procedures neglect possible developmental delays between subjects, or some pathology affecting the cardiac pace, for instance, a key feature that we precisely want to detect. A spatiotemporal registration scheme should register individual growth scenarios both in space (usual geometrical variations of the anatomy) and in time (change of the speed of evolution). Time changes should put the ages of the subjects into correspondence, which represents the same developmental stage.

In Declerck et al. (1998), a deformation of the 4D domain is provided via 4D planispheric transformations for the registration of the heart motion. In Perperidis et al. (2005), spatiotemporal deformations are computed. The temporal part is a 1D function showing the change of cardiac dynamics between the source and the target subject. This temporal alignment is performed jointly with the registration of the anatomy. These methods focus on the registration between a *pair* of individual trajectories, and requires a fine temporal sampling of the trajectories. A main contribution of our paper will be to use such spatiotemporal deformations for the inference of statistical properties at the population level, via the estimation of spatiotemporal atlases.

2.1.4 Ingredients for a Spatiotemporal Statistical Model

This review of the literature shows that several aspects of the design of a 4D statistical analysis have been addressed separately by different authors, in different contexts and with different tools. There is a lack of a consistent framework to embed these concepts together, covering the estimation of individual trajectories and the inference of population statistics.

In light of this review, a statistical framework for longitudinal data analysis might include:

- The estimation of a *continuous* shape evolution from a set of observations sparsely distributed in time. These individual trajectories could be used to compare the anatomy of two subjects, who have not been scanned at the same age. They could also be used to analyze the speed of evolution of a given subject at any time-point.
- The comparison between individual trajectories, which should measure not only morphological differences (commonly described by 3D deformations) but also the temporal re-alignment which put the developmental stages of different subject into correspondence. This temporal re-alignment will detect different speeds of evolution and therefore possible developmental delays between subjects.
- A generative statistical model, which combines the two previous concepts to estimate evolution patterns that are shared among a given population. The estimated statistics should include a mean (a growth scenario representative of the population) and variance (the typical variations of this mean growth scenario evident in the population).

2.1.5 Terminology

The survey of the current literature also raises the problem of terminology: there is no consensus among authors about which words refer to which concepts. In this paper, we will use the following definitions:

Data:

- **cross-sectional data** is a set of samples, which are supposed to be comparable, or homologous (like samples drawn from a healthy adult population, for instance). No notion of time is involved, or equivalently, the effect of time or age on the data can be neglected.
- **time-series data** is a set of data that are indexed by any temporal marker like age, indicator of developmental stage, disease progression or index of a frame in a movie, for instance. No assumption is made that a sub-set of the samples correspond to the same object seen at different time-points.
- **longitudinal data** is a time-series data set, which contains repeated observations of individual subjects over a period of time. As a consequence, each subject in the data set should have been observed more than once at different time-points.

Methods:

- **shape/image regression**, also called tracking, refers to the estimation of a continuous evolution model from a time-series data set. This tool estimates shape changes between discrete temporal observations or averages time-indexed observations into a single evolution.
- **spatiotemporal registration** puts two individual trajectories into correspondence. This involves the notion of correspondence between shapes and between time-points.
- **spatiotemporal or longitudinal data analysis** measures the similarities and the differences between individual trajectories. It takes into account the fact that individual subjects were observed several times, which makes it

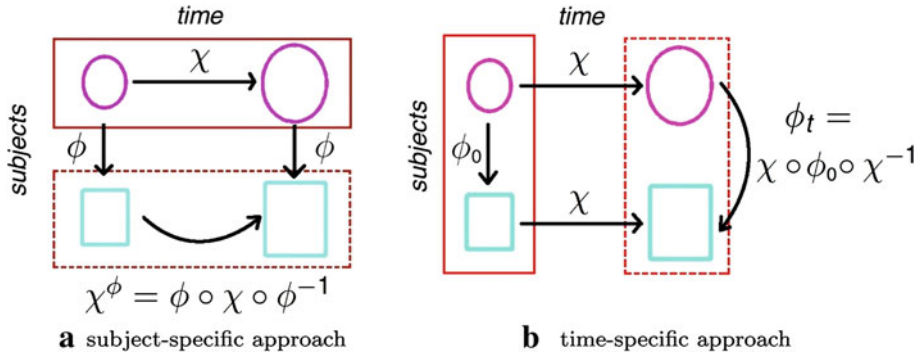


Fig. 2 Illustration of the hypotheses underlying the subject- and time-specific approaches. In the subject-specific approach (*left*), one considers that one subject is "*circle*" and the other is "*square*": the difference is described by a single function ϕ, which maps *circles* to *squares*. The evolution of the first subject is described by a function χ which maps a *small circle* to *big circle*. As a consequence, the evolution of the second subject is described by another function χ^{ϕ} which maps a *small square* to a *big square*. In the time-specific approach (*right*), one describes the evolution by a universal function χ, which tends to scale the shapes. At the first-time point, the difference between subjects is described by a function ϕ which maps the *small circle* to the *small square*. At a later time, the inter-subject variability has changed according to χ: now the difference between subjects is described by ϕ_t which maps a *big circle* to a *big square*

more constrained than the analysis of the effect of time on the observations.

According to these definitions, shape or image regression may be performed on time-series data, whereas spatiotemporal analysis can only be performed on longitudinal data.

2.2 Two Possible Generative Models for Longitudinal Data

2.2.1 Spatiotemporal Variations of a Typical Growth Model

A generative statistical model is a set of hypotheses, which explain how individual trajectories could be derived one from the others. In other words, it should provide an answer to the two fundamental questions: given the anatomy of one subject at time t,

- how can we predict the anatomy of this subject at a later time $t' > t$?
- how can we derive the typical anatomy of another subject at the same time-point?

Once these answers are provided, we can easily define a generative statistical model at the population level. This model will assume the existence of a mean growth scenario representative of the population, such that the individual trajectories can be seen as a derivation of this mean scenario. The mean scenario captures the invariants in the population and detects the growth patterns, which are shared among the subjects. The derivation of the mean scenario captures the variance of this mean configuration within the population.

In light of the literature survey, there are at least two different ways to answer these questions. We refer to these two paradigms as a "subject-specific approach" and a "time-specific approach".

2.2.2 Subject-Specific Approach

In the subject-specific approach, a specific reference frame is attached to each subject. The *whole* evolution of each subject is described within the same reference frame: the reference frames are atemporal. We assume that there is a template reference frame in which the evolution is written by a time-varying shape $M(t)$: the prototype scenario of evolution in the population, which can be seen as the $4D$ analog to the template shape in $3D$. We usually assume that the time-varying shape derives continuously from a template shape M_0 at a reference time point t_0. This is written as: $M(t) = \chi_t(M_0)$, where χ_t is a smoothly varying 3D deformation called the growth function ($\chi_{t_0} = \mathrm{id}$ so that $M(t_0) = M_0$).

Change of coordinates from the template reference frame to each subject's reference frame is modeled by 3D deformations ϕ^s. Since these reference frames are atemporal, the deformations ϕ^s do *not* depend on time. As a consequence, the evolution function $M(t)$ has a different expression in each coordinate system: the evolution of a subject S is given as $S(t) = \phi(M(t))$, also written as $S(t) = \phi(\chi_t(M_0))$. It is as if a single object ($M(t)$) is seen by different observers in different coordinate systems. As illustrated in Fig. 2a, the change of coordinates ϕ^s transports the evolution function $\chi(t)$ from the template frame to the subject's frame: $\chi_t^S = \phi^s \circ \chi_t$, so that $S(t) = \chi_t^s(M_0)$. The evolution mapping is therefore *specific* to each subject.

In this modeling, we can include time as an additional variable, so that the reference frame of each subject is described by 3 spatial coordinates and 1 temporal coordinate. This

means that both the anatomy and the age is relative to the subject. This specific time variable can be called the "physiological age" of the subject, as if each subject has their own biological clock. In the reference frame of the prototype, the time would be the absolute age, computed from the date of birth. Then, the $3D$ warp $\phi(x, y, z)$ needs to be generalized to a deformation of the underlying $4D$ space: $\Phi(x, y, z, t)$.

The most general form of a $4D$-deformation is $\Phi(x, y, z, t) = (\phi(x, y, z, t), \psi(x, y, z, t))$, where $\phi(x, y, z, t)$ denotes the 3 spatial coordinates of $\Phi(x, y, z, t)$ (the morphological deformation) and $\psi(x, y, z, t)$ its temporal coordinate (the time warp).

Assuming that $\psi(x, y, z, t)$ depends on the spatial variables (x, y, z) means that different parts of the anatomy of a given subject would evolve at different speeds. This is definitely possible in applications involving multi-shape comparisons. However, in this paper, we will assume that all points of the anatomy of a given subject have always the same physiological age over time. In this case, ψ depends only on the time variable t: $\psi(t)$. This assumption is likely to be valid in most longitudinal studies, focusing on one specific structure. The time warp $\psi(t)$ maps the absolute age in the reference frame of the prototype to the physiological age of a given subject. Note that this function should be monotonic, assuming that the sequence of events in every individual trajectory occur in the same order (from birth onwards) but at a different pace.

Since the change of coordinate maps ϕ are independent of time, they are of the form: $\phi(x, y, z)$. Therefore, in the subject-specific setting, the $4D$ deformations are written as:

$$\Phi(x, y, z, t) = (\phi(x, y, z), \psi(t)).$$

The morphological deformation ϕ is used to measure the geometrical variability. The time warp ψ is used to detect possible developmental delays between subjects.

Note that the most general form of $4D$ deformations, without any assumptions on the temporal dependency of the spatial part $\phi(x, y, z, t)$, cannot be used in a statistical model. Indeed, such models will be not identifiable, as there would be an infinite number of different spatial/temporal combinations to explain the same data set.

2.2.3 Time-Specific Approach

In the time-specific approach, every subject is embedded into the same reference frame, which transports everyone over time. It is as if different objects are seen by a single observer. More precisely, there is a common reference frame at reference time $t = 0$ (the "origin of the world") in which the anatomy of every subject is described. The evolution function χ_t changes the geometry of this reference frame over time. At each time t, there is one single reference frame which embeds the anatomy of every subject: this frame is universal. The same function χ_t applies for each subject, so that the evolution of any subject is given by $S(t) = \chi_t(S_0)$, where S_0 represents the anatomy of the subject at the reference time t_0.

In the common reference frame at $t = 0$, we assume that each subject's anatomy S_0 results from a deformation of the prototype anatomy M_0: $S_0 = \phi^s(M_0)$. The deformations ϕ^s describe the inter-subject variability at time $t = 0$. In this framework, the mapping between the template and the subject shape changes over time according to the evolution function χ_t. At a later time t, the template has evolved as $M(t) = \chi_t(M_0)$ and the subject shape has evolved as $S(t) = \chi_t(S_0)$. This shows that the template-to-subject registration has become: $S(t) = \phi_t(M(t))$ where $\phi_t = \chi_t \circ \phi \circ \chi_t^{-1}$, as illustrated in Fig. 2b. Whereas the evolution function is independent of the subject, the inter-subject variability is specific to time.

We can also include possible developmental delays in this framework. If χ_t is a universal function which carries the anatomies over time, we can imagine that every subject follow this universal scenario at its own pace. There is a subject specific time warp ψ, so that the evolution of this subject is given by $\chi_{\psi(t)}$. However, we must admit that this time-realignment fits less naturally into this time-specific framework than for the subject-specific framework. In particular, it is not clear how to distinguish a developmental delay from a variation of the inter-subject variability in this setting.

This time-specific approach also defines a deformation of the underlying $4D$-space $\Phi(x, y, z, t)$. The morphological deformations ϕ now depend on time according to the evolution function χ_t. This leads to the particular form of the $4D$ mapping:

$$\Phi(x, y, z, t) = (\phi(x, y, z, \psi(t)), \psi(t)),$$

where the geometrical part has the form:

$$\phi(x, y, z, t) = \chi_t \circ \phi_0 \circ \chi_t^{-1}(x, y, z).$$

This last equation is the constraint, which eventually makes the statistical model identifiable.

2.2.4 Which Method for Which Problem?

The subject-specific approach focuses on the variations of a growth scenario from subjects to subjects. One is interested in analyzing how individual trajectories vary across subjects. The time-specific approach focuses on the evolution of the inter-subject variability over time. One is more interested in the evolution of the statistical properties (mean and variance) of the population over time, as illustrated in Fig. 3. These two approaches are based on different assumptions and lead to different statistical estimations.

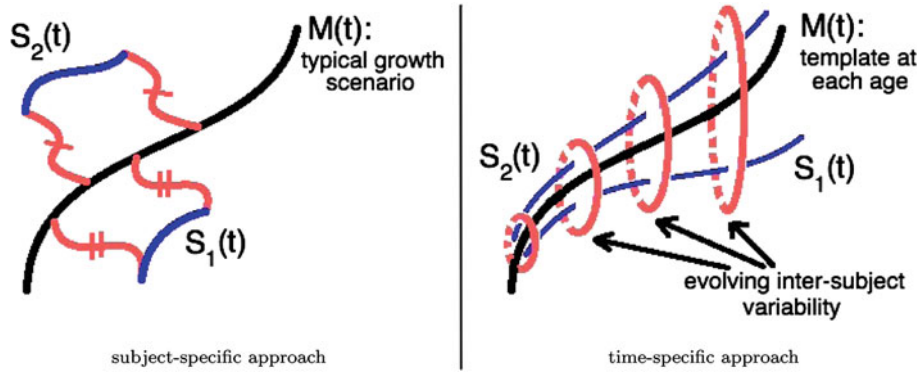


Fig. 3 Subject- versus time-specific approach. In the subject-specific approach (*left*) the mean scenario averages the individual trajectories. The inter-subject variability is supposed to be constant over time. In the time-specific approach (*right*), every subject is supposed to follow the same mean scenario of evolution, up to a change of the initial conditions. The mean scenario describes how the inter-subject variability evolves over time

The subject-specific approach is the only one to take into account change of coordinates between subjects, and therefore the only one to accommodate for scaling effects across subjects. The time-specific paradigm uses a single diffeomorphic deformation to describe the evolution of every subject. This assumes that the structure of two different subjects, which are superimposed at one time, will remain superimposed in the future. Such topological constraints are often unrealistic. Moreover, the statistical estimations in the subject-specific paradigm are more robust when the number of subjects is greater than the number of observations per subjects, which is the case with the longitudinal data set on which we aim at applying this methodology in Sects. 4 and 5. For these reasons, the presented work will focus on the subject-specific paradigm.

3 A Subject-Specific Approach using 3*D* Diffeomorphisms and 1*D* Time Warps

In this section, we propose an instance of the subject-specific paradigm for the analysis of longitudinal shape data, given as point sets, curves or surfaces. Among several other possible choices, we will build our methodology on the large diffeomorphic deformations setting for defining the registration between shapes. This setting is particularly adapted to define statistical models using deformations due to the metric properties of the considered space of diffeomorphisms (Vaillant et al. 2004; Durrleman et al. 2009a). In particular, we will propose an extension of this framework to construct monotonic 1D functions for our "time warp" in a very generic way. We will also consider the geometrical shape like curves and surfaces as currents (Glaunès 2005). This allows us to inherit from the statistical and computational tools introduced in Durrleman et al. (2009a) and Durrleman (2010) for the estimation of representative shapes, called templates.

We follow the approach in three steps outlined in the Introduction: (i) the estimation of individual growth trajectories via the inference of a growth model, (ii) the comparison of individual trajectories based on a morphological map and a time warp, (iii) the estimation of statistics from a set of individual trajectories: mean scenario of evolution and analysis of the spatiotemporal variability.

3.1 Sketch of the Method

3.1.1 Growth Model for Individual Shape Evolution

Our purpose is to fit a *continuous* shape evolution to a *discrete* set of shapes (S_i) of the same subject acquired at different time points (t_i). To infer such a continuous shape evolution, we need a prior on the growth of the shape, called a "growth model". Here, we hypothesize that the baseline shape S_0 observed at time $t = 0$ continuously and smoothly deforms over time. To be more precise, our growth model assumes that the evolution of the shape S_0 can be described by a continuous flow of diffeomorphisms χ_t. This means that for each t varying in the interval of interest $[0, T]$, χ_t is a diffeomorphism of the underlying 2D or 3D space, which models the smooth and invertible deformation which maps the baseline at $t = 0$ to its actual shape at time t. The diffeomorphisms vary continuously over time (the deformation $\chi_{t+\delta t}$ is close to χ_t). Mapping the baseline shape S_0 with the time-varying functions χ_t leads to a continuously deforming shape $S(t) = \chi_t(S_0)$: the individual trajectory of the considered subject. Note that this imposes that $\chi_0 = \mathrm{Id}$, the identity map, so that $S(0) = S_0$.

Given the set of discrete observations (S_i) at time-points t_i, one needs to estimate the flow of diffeomorphisms χ_t, which

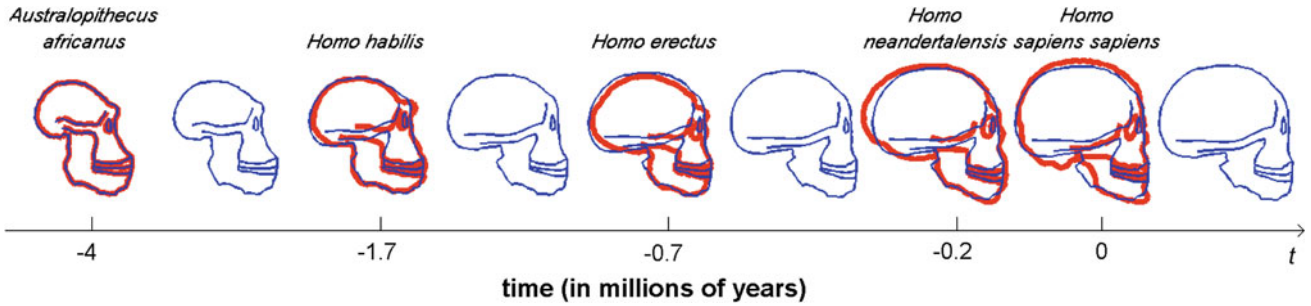


Fig. 4 Shape regression of a set of five $2D$ profiles of hominid skulls (in *red*). The *Australopithecus* profile is chosen as the baseline S_0. The temporal regression computes a *continuous* flow of shapes $S(t)$ (here in *blue*) such that the deforming shape matches the observations at the corresponding time-points. It is estimated by fitting a growth model, which assumes a diffeomorphic correspondence between the baseline and every stage of evolution ($S(t) = \chi_t(S_0)$), with the diffeomorphism χ_t varying continuously in time

may have led to these observations. A Maximum A Posteriori (MAP) estimation, in the same framework as in Durrleman (2010, Chap. 5) leads the minimization of the discrepancy between the growth model at time t_i ($S(t_i) = \chi_{t_i}(S_0)$) at the actual observation S_i, up to a regularity constraint on the smoothness of the flow of diffeomorphisms $(\chi_t)_{t\in[0,T]}$:

$$E(\chi) = \sum_{t_i} d(\chi_{t_i}(M_0), S_i)^2 + \gamma^{\chi}\mathrm{Reg}(\chi) \tag{1}$$

where d is a similarity measure between shapes, which will be the distance on currents in the following, $\mathrm{Reg}(\chi)$ a regularity term, which will be the total kinetic energy of the deformation, and γ^{χ} a scalar parameter quantifying the trade-off between regularity and fidelity to data. The optimization of this criterion will be explained in Sect. 3.2.2.

As an illustrative example, we used five $2D$ profiles of hominid skulls which consist of six lines each,[1] as shown in Fig. 4. Each profile correspond to a hominid (*Australopithecus*, *Homo habilis*, *Homo erectus*, *Homo Neandertalensis* and *Homo sapiens sapiens*) and is associated to an age (in millions of years). The regression infers a continuous evolution from the *Australopithecus* to the *H. sapiens sapiens* which matches the intermediate stages of evolution.

If there is only one data S_1 at time $t_1 = T$, the criterion (1) defines the registration of S_0 to S_1. In the LDDMM framework, the result of such a registration is a geodesic flow of diffeomorphism between $t = 0$ and $t = T$ that maps S_0 close to S_1 (Miller et al. 2002). With several data at successive time points, we will show in Sect. 3.2.2 that the result is a flow of diffeomorphism which is geodesic only between successive time points (i.e. piecewise geodesic). We will also show that the computation of the regression functions χ_t takes into account all the observations S_i in the past and future simultaneously. Therefore, it differs from pairwise registration between consecutive shapes. For instance, if the trade-off γ^{χ} tends to infinity (no fidelity-to-data term) the regression is a constant map $\chi_t = \mathrm{Id}$ for all t. As γ^{χ} decreases, the piecewise geodesic regression matches the data with increasing "goodness of fit". This framework allows us also to perform the regression even if several data are associated to the same time-point. This will be used in Sect. 5 to estimate a mean growth scenario of a time-series cross-sectional data set.

Note that if T is greater than the latest time-point of the data $t_{\max}$, then the regression function χ is constant over the interval: $[t_{\max}, T]$. Therefore, the method extrapolates with constant shape outside the time interval $[0, t_{\max}]$. Such an extrapolation will be needed to compare the evolution of two subjects, whose latest observation correspond to different time-points. Similarly, we can also extrapolate the evolution function at time earlier than 0 with a constant map, so that the evolution function can be defined on any arbitrary time interval.

3.1.2 Spatiotemporal Registration Between Pairs of Growth Scenarios

We suppose now that we have two subjects S and U which have been scanned several times each (but not necessarily the same number of times and possibly at different ages). Let S_{t_i} (resp. U_{t_j}) be the shapes of subject S (resp. U) at ages t_i (resp. t_j). We define a time-interval of interest which contains every t_i's and t_j's. Without loss of generality, we can assume that this time interval of interest is of the form $[0, T]$.

We infer an individual growth model $S(t)$ from the data of the source subject $\{S_{t_i}\}$, using the procedure of the previous section. As a result, the continuous shape evolution $S(t)$ is of the form: $S(t) = \chi_t(S_0)$ for $t \in [0, 1]$.

Our goal is to define a spatiotemporal deformation of the continuous evolution $S(t)$ into $S'(t)$ so that the deformed shapes $S'(t_j)$ at the time-points of the target t_j match the shape U_{t_j} (thanks to the continuous regression, we can define

[1] Source: www.bordalierinstitute.com.

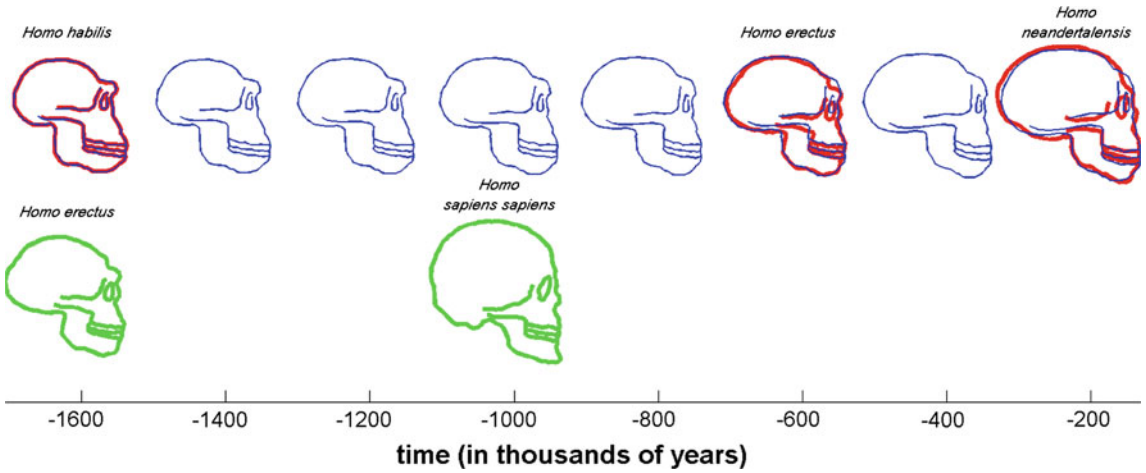


Fig. 5 Illustrative pairwise registration: data preparation. The database is cut in two to compare the evolution {*H. habilis-erectus-neandertalensis*} (*red shapes*) to the evolution {*H. erectus-sapiens sapiens*} (*green shapes*). The later evolution is translated in time, so that both evolutions start at the same time. Then, one performs a shape regression of the source shapes (*blue shapes*). The spatiotemporal registration of this continuous source evolution to the target shapes is shown in Figs. 6 and 7

$S(t_j)$ for the target time point t_j even if the source has not been observed at this age.) For this purpose, we introduce two functions (using the subject-specific paradigm in Sect. 2): the 3D morphological deformation ϕ and the 1D time warp ψ. The morphological deformation is a 3D-function, which maps the geometry of the source to the geometry of the target (change of reference frame). Every frame of the source sequence $S(t)$ is deformed using the same function. The time warp ψ maps the time-points t within the time interval $[0, T]$ to $\psi(t)$. This function does not change the frames of the sequence $S(t)$ but change the speed at which the frames are displayed. It models the change of the dynamics of the evolution of the source with respect to the evolution of the target. We impose this 1D function to be monotonic, assuming that the shape changes occur in the same order, even if at a different pace between source and target. The combination of these two functions gives the spatiotemporal deformation of the continuous evolution $S(t)$, defined as:

$$S'(t) = \phi\left(S\left(\psi(t)\right)\right). \tag{2}$$

Using the fact that $S(t) = \chi_t(S_0)$, this becomes:[2]

$$S'(t) = \phi(\chi_{\psi(t)}(S_0)). \tag{3}$$

In a MAP setting, the estimation of the best possible spatiotemporal deformation (ϕ, ψ) of the source evolution which fits the the target observations, leads to the minimization of the discrepancy between the deformed source at target's time-points $S'(t_j) = \phi(S(\psi(t_j))$ and the target's shape U_{t_j}:

$$E(\phi, \psi) = \sum_{t_j} d\left(\phi(S(\psi(t_j))), U_{t_j}\right)^2 + \gamma^{\phi}\mathrm{Reg}(\phi) + \gamma^{\psi}\mathrm{Reg}(\psi), \tag{4}$$

where d is a distance between shapes, $\mathrm{Reg}(\phi)$ and $\mathrm{Reg}(\psi)$ the measure of regularity of the deformation ϕ and ψ and $\gamma^{\phi}, \gamma^{\psi}$ the usual scalar trade-offs between regularity and fidelity to data.

An illustration of spatiotemporal registration is shown in Figs. 5, 6 and 7. We use the same set of profiles of $2D$ hominids skulls as in Fig. 4. Here we want to compare the evolution {*H. habilis-Homo erectus-H. neandertalensis*} (called earlier evolution) with the evolution {*H. erectus-H. sapiens sapiens*} (called later evolution). The differences may be due to a change of the shape of the skull, as well as a change of the dynamics of evolution between the earlier and the later evolution. Therefore, we divide the database into two groups, considered as two different subjects, and translate the target back four million years, so that both evolutions start at the same time (this can be seen as a "rigid" temporal alignment as a pre-processing). See Fig. 5.

The regression of the source data leads to a continuous source evolution $S(t)$ shown in blue in the first row of Fig. 5. The estimation of the spatiotemporal deformation between the source and the target results in a morphological deformation ϕ and a time warp ψ, see Fig. 6. The morphological deformation shows that the jaw is less prominent and the skull larger and rounder during the later evolution than during the earlier evolution (second row in Fig. 6). The effect of the time warp is to accelerate the source evolution to adjust to the rate of shape change between the target shapes (third row in Fig. 6). The graph of the time warp is plotted in

[2] We notice that in the time-specific paradigm, this would be $S''(t) = \chi_{\psi(t)}(\phi(S_0))$ (see Sect. 2).

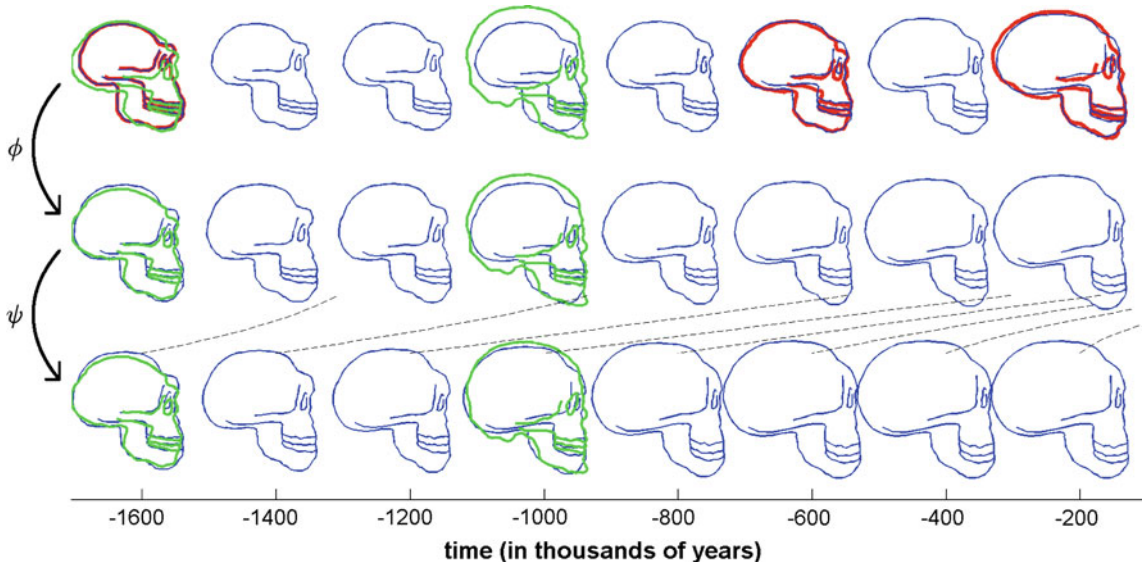


Fig. 6 Illustrative pairwise registration: morphological deformation and time warp. *Top row* The input data as prepared in Fig. 5 with the continuous source evolution (*blue*) superimposed with the target shapes (*green*). *Middle row* The morphological deformation ϕ is applied to each frame of the source evolution. It shows that, independently of time, the skull is larger, rounder and the jaw less prominent during the later evolution relative to the earlier evolution. *Bottom row* The time warp ψ is applied to the evolution of the second row. The *blue shapes* are moved along the time axis (as shown by *dashed black lines*), but they are not deformed. This change of the speed of evolution shows an acceleration of the later evolution relative to the earlier evolution. Taking this time warp into account enables a better alignment of the source to the target shapes than only the morphological deformation. Note that the morphological deformation and the time warp are estimated simultaneously, as the minimizers of a combined cost function

Fig. 7a. It shows an almost linear increase in speed. The slope of the curve is of 1.66, thus meaning that the later evolution evolves 1.66 times faster than the earlier evolution. This value is compatible with the growth speed of the skull during this period according to the values reported in the literature and in Fig. 7b: between *Homo erectus* and *H. sapiens sapiens* the skull volume had grown at a rate of $(1500-900)/0.7 = 860\,\mathrm{cm}^3$ per millions of years, whereas between *H. habilis* and *Homo neandertalensis*, it had grown at $(1500-600)/1.7 = 530\,\mathrm{cm}^3$ per millions of years, namely 1.62 times faster.

3.1.3 Atlas Estimation from Longitudinal Data Sets

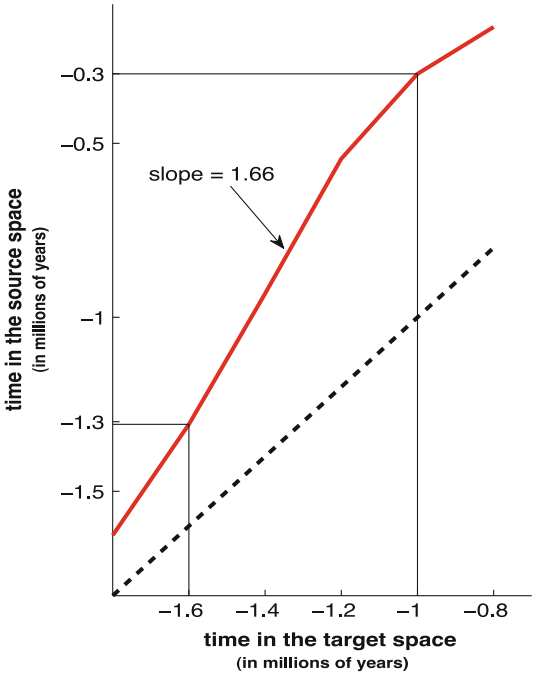
In this section, we want to combine the previously introduced growth model and spatiotemporal deformations to estimate statistics from a longitudinal database. Given the repeated observations of a group of subjects, we assume that each subject's evolution derives from the same prototype evolution, called a "mean scenario of evolution". Each subject-specific evolution is derived from the mean scenario via its own spatiotemporal deformation. The analysis of the set of all the spatiotemporal deformations in the population will lead to the estimation of the typical variations of the mean scenario in the population (the variance of the population in a sense to be defined). We assume that the mean scenario of evolution is given by the growth model of an unknown prototype shape M_0, called template in the sequel.

Formally, this means that there is a growth function χ_t for $t \in [0, T]$ and a template shape M_0, so that the mean scenario of evolution is written as: $M(t) = \chi_t(M_0)$ with $M(0) = M_0$. For each subject s ($s = 1, \ldots, N_{\mathrm{subj}}$), the subject-specific spatiotemporal deformation of the mean scenario is written as: $S^s(t) = \phi^s\,(M\,(\psi^s(t)))$ for all $t \in [0, T]$. ϕ^s is the morphological deformation for subject s and ψ^s its time warp. These two functions model how the anatomy of the subject and the dynamics of evolution can be derived from the prototype scenario of evolution. Eventually, we suppose that the observation of the subject s at time-point t_j^s, denoted S_j^s, is the temporal sample from $S^s(t)$ at time point t_j^s, up to a random Gaussian noise: $S_j^s = S^s(t_j^s) + \varepsilon_j^s$:

$$S_j^s = \phi^s\left(\chi_{\psi^s(t_j^s)}(M_0)\right) + \varepsilon_j^s, \tag{5}$$

where the Gaussian variables ε_j^s are independent and identically distributed over the subject-index s and the time-index j.

This equation is our generative statistical model, which explains how the observations can be seen as instances of a random process. The fixed parameters are the prototype shape M_0 and the growth function χ_t. The random parameters are the spatiotemporal deformations (ϕ, ψ) (each estimated (ϕ^s, ψ^s) is an instance of these random deformations). Both



a time warp



b skull volume evolution

Fig. 7 Illustrative pairwise registration: analysis of the time warp. *Top* plot of the 1D time warp $\psi(t)$ putting into correspondence the time-points of the target shapes with that of the source. The $x = y$ line (*dashed in black*) would correspond to no dynamical change between source and target ($\psi(t) = t$). The *slope* indicates that the shape changes between target data occur 1.66 times faster than the changes in the source evolution, once morphological differences has been discarded. *Right* the graph of the skull volume over the human evolution as found in the literature (*source* www.bordalierinstitute.com). This *curve* shows that the increase in skull volume between *H. erectus* and *H. sapiens sapiens* was 1.62 times faster than between *H. habilis* and *H. neandertalensis* (ratio between the slope of the two straight lines). This value is compatible with the acceleration measured by the time warp: 1.66

the fixed and the random parameters are unknown and should be estimated given the actual observations.

In the same MAP setting as in the previous section, the estimation of the unknown parameters can be done by minimizing the following combined cost function:

$$E\left((\psi^s)_{s=1,\ldots,N_{\text{subj}}}, (\phi^s)_{s=1,\ldots,N_{\text{subj}}}, \chi, M_0\right) = \sum_{s=1}^{N_{\text{subj}}} \left\{ \sum_{t_j^s} \mathrm{d}(\phi^s(\chi_{\psi^s(t_j^s)} M_0), S^s(t_j^s))^2 + \gamma^{\phi}\mathrm{Reg}(\phi^s) + \gamma^{\psi}\mathrm{Reg}(\psi^s) + \gamma^{\chi}\mathrm{Reg}(\chi) \right\} \quad (6)$$

The output is the prototype shape M_0, the growth function χ_t and the set of spatiotemporal deformations ϕ^s, ψ^s for every subject s. These variables are called a "spatiotemporal atlas". In Sect. 3.3, we will show how we can perform statistics on the estimated deformations (ϕ^s, ψ^s), like Principal Component Analysis for instance. Such statistics will describe the changes in shape and the variations of the speed of evolution across subjects.

To illustrate the method, we run the atlas estimation given the two "subjects" in Fig. 5: the first subject (in red) consists of three shapes, the second subject (in green) consists of two shapes. From these five shapes, the method returns the estimated template, the mean scenario and the two spatiotemporal registrations of this mean scenario to each subject. The estimated template M_0 is given as a current, which does not form a set of curves anymore (Durrleman et al. 2009a). To give an illustration of the atlas, we map the youngest shape of each subject to this current and pick the deformed shape that is the closest to the estimated template. Then, one runs one more iteration of the atlas algorithm, to show the mean scenario and the spatiotemporal registrations as deformations of this template shape. This is shown in Fig. 8. In particular, the two time warps which put into correspondence the evolution stages of each subject to the ones of the estimated mean scenario are shown in Fig. 9b.

Let us denote $S_1(t)$ the spatiotemporal deformation of the mean scenario, which is supposed to match the shape of the first subject: $S_1(t) = \phi_1(M(\psi_1(t)))$. Similarly, $S_2(t) = \phi_2(M(\psi_2(t)))$ matches the shapes of the second subject. Then, by definition, we have: $S_2(t) = \phi_2\left(\phi_1^{-1}(S_1(\psi_1^{-1}(\psi_2(t))))\right)$. At the first glance, this suggests that $(\phi_2 \circ \phi_1^{-1}, \psi_1^{-1}\psi_2)$ corresponds to the spatiotemporal registration between the first subject (considered then as the source) to the second subject (considered as the target). We superimposed in Fig. 9c the graph of $\psi_1^{-1} \circ \psi_2$ with the graph of the time warp estimated in the previous section and shown in Fig. 7a. As expected, the two curves show a similar pattern, namely the overall acceleration of the source relative to the target. However, noticeable differences appear, in particular

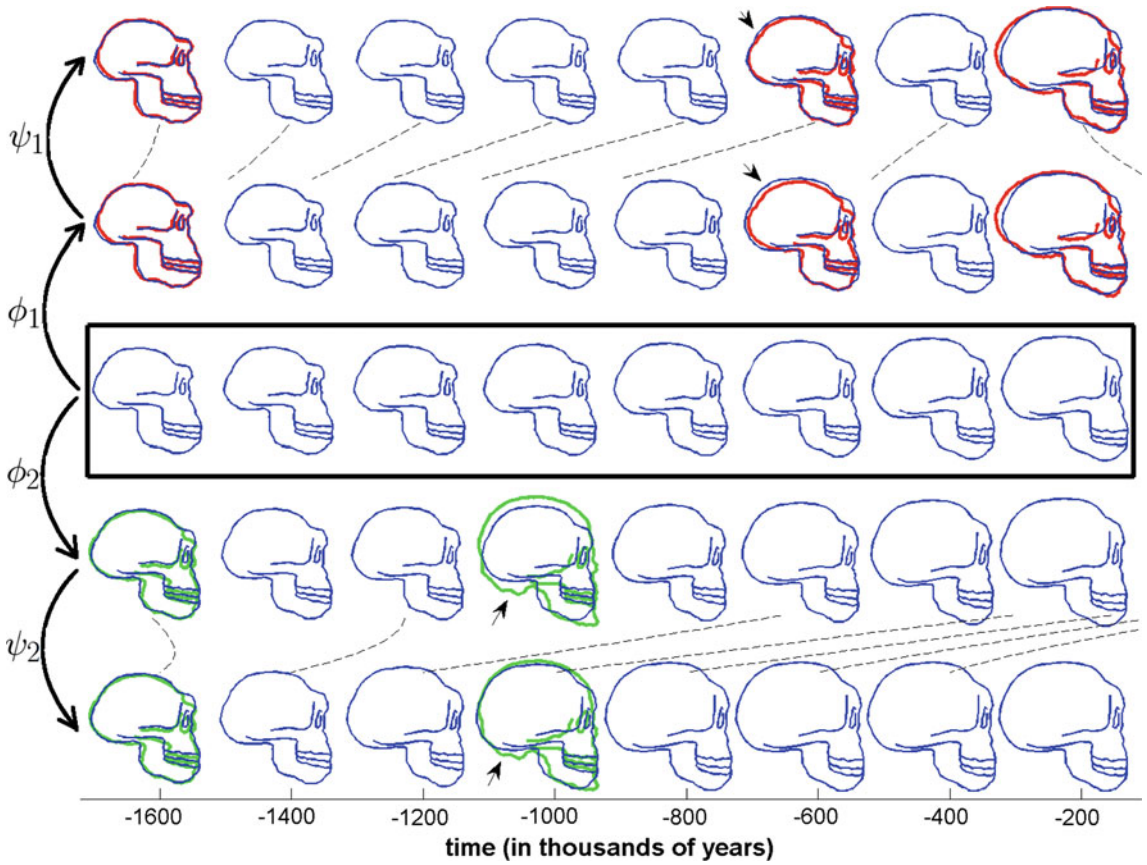


Fig. 8 Spatiotemporal atlas estimation given the two "subjects" in Fig. 5. On the *middle row* is shown the estimated mean scenario of evolution. The first frame of this scenario (*far left*) is the estimated template shape. Two *upper rows* represents the morphological deformation and then the time warp, which jointly maps the mean scenario to the shapes of the first subject (*red shapes*). Tow *lower rows* represents the spatiotemporal deformation of the mean scenario to the shapes of the second subject (*green shapes*). *Black arrows* indicate areas where the most important shape deformations occur

in the slope of the curves. This can be explained by at least two reasons. First, what we called here $S_1(t)$ is *not* the same shape evolution as the one computed in the pairwise registration case (first row in Fig. 5): the regression of the source subject in the registration case did not take into account any information about the target shapes, whereas the mean scenario $M(t)$ (and consequently its deformation $S_1(t)$) averages the growth patterns of *both* subjects. Second, the reasoning above does not take into account the residual errors into account: assuming that $\tilde{S}_1(t)$ and $\tilde{S}_2(t)$ are the true evolution of each subject, then we have: $\tilde{S}_1(t) = \phi_1(M(\psi_1(t))) + \varepsilon_1(t)$ and $\tilde{S}_2(t) = \phi_2(M(\psi_2(t))) + \varepsilon_2(t)$, where $\varepsilon_1(t)$ and $\varepsilon_2(t)$ models the residuals shape which contains noise, small-scale variations and everything else, which cannot be explained by the model. The squared norm in the criterion to be optimized shows that we assume these residuals to be Gaussian random variables (see Durrleman 2010 for more details). This shows therefore that $\tilde{S}_2(t) = \phi_2 \circ \phi_1^{-1}(\tilde{S}_1(\psi_1^{-1} \circ \psi_2)(t) + \phi_2 \circ \phi_1(\varepsilon(\psi_1^{-1} \circ \psi_2(t)))$, meaning that the residual error between the two scenarios $\tilde{S}_1(t)$ and $\tilde{S}_2(t)$ is no more Gaussian. Therefore, to retrieve the same deformations and time warps, one would need to change the squared norm in the registration criterion to take into account the distortion in the distribution of the residuals induced by the deformations.

The discussion above highlights the main features of the atlas construction method. The main assumption is that the different subjects derive from the same prototype scenario, and therefore share common growth patterns even if altered in their shape and timing. The atlas aims precisely at detecting these common features and the variations of their shape and pace in the population. Every pattern, which is specific to a given individual, is discarded from the atlas and remains in the residuals. In this sense, the atlas is a statistical tool, which detects the reproducible patterns in the population. Compared to pairwise registration, the advantages of the atlas construction is that it can be applied to more than two subjects and that it does not favor any particular subject in the population.

Remark 1 (On the assumption of diffeomorphic maps) In this modeling, we suppose that the evolution function χ and the

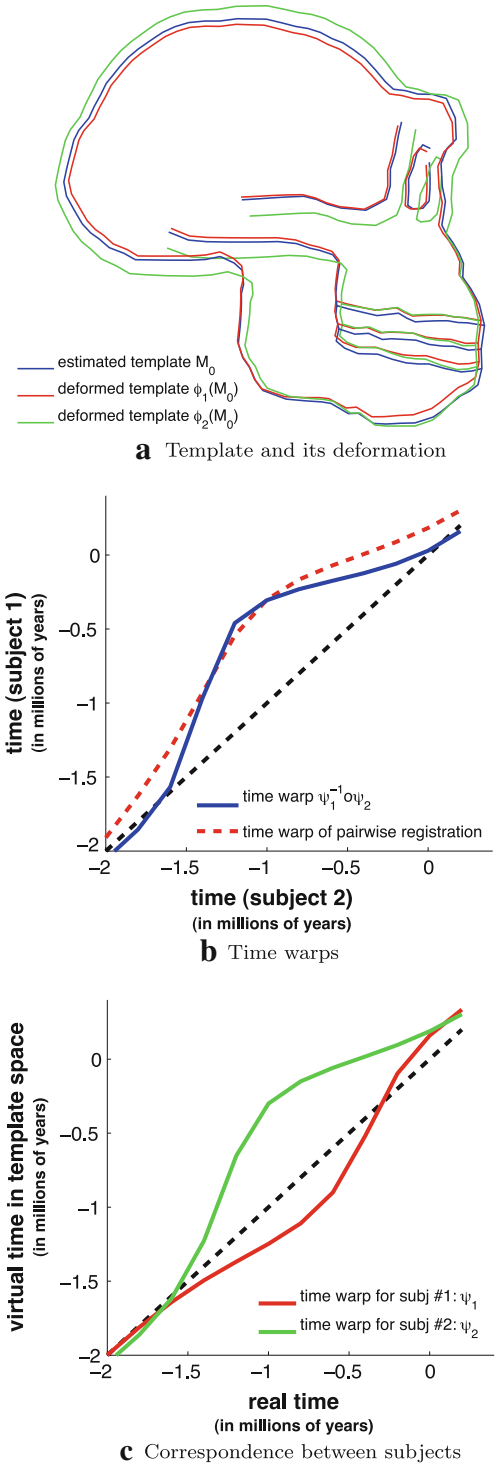


Fig. 9 Spatiotemporal atlas estimation: template and time warp. (**a**) The template and its morphological deformation to the first subject (*red*) and the second subject (*green*). This corresponds to the far left frames in the second, third and fourth row in Fig. 8. (**b**) The graphs of the two time warps, mapping the subjects' growth speed to that of the mean scenario. When the curve is above $x = y$ axis, the subject's evolution is in advance relative to the rate of shape changes given by the mean scenario. (**c**) Graph of the function: $\psi_1^{-1} \circ \psi_2$ (in *blue*), which maps the dynamics of the two subjects. The *dashed red curve* is the time warp given by the pairwise registration as shown in Fig. 7a

morphological deformations ϕ are 3D diffeomorphisms and that the time warps ψ are 1D diffeomorphisms. The motivation and consequences of choosing diffeomorphic maps are different in each case.

The evolution function χ maps the anatomy of a subject over time. Setting χ as a diffeomorphism assumes a smooth one-to-one correspondence between any observed shapes of the same subject. This includes modes of growth like atrophy, dilatation, torque, etc. However, this cannot model a tearing of the shape, its division or the creation of another disconnected component over time. This assumption is realistic in many practical case, like for the heart over a cycle or the macroscopic observation of a brain structure during infancy.

The morphological deformations ϕ model the geometrical inter-subject variability. Assuming a smooth one-to-one correspondence between the anatomies of two different subjects is more questionable. As highlighted in Durrleman et al. (2011) and Durrleman (2010, Chap. 5), the diffeomorphism is used to decompose the inter-subject variability into two terms: the diffeomorphic geometrical variability captured in the deformations and the non-diffeomorphic variability in terms of "texture" captured in the residuals (modeled by the random Gaussian variables). Both terms can be used for the statistical analysis, whereas in this work we will focus only on the geometric variability captured by the deformations. Extending the work of Durrleman et al. (2011) to analyze the non-diffeomorphic variations would be possible but out of the scope of this paper.

The time warp ψ model the change of speed of evolution between subjects. The diffeomorphic assumption in 1D implies that the function is smooth and monotonic. The monotonic property assumes that the sequence of the events during evolution occur in the same order for every subject (from birth to death). This is a very realistic (if not desirable) hypothesis, at least from a biological point of view. Moreover, assuming the evolution of a structure is smooth (at least differentiable) like its inverse is also very realistic, so that one can speak about the *speed* of an evolution. Therefore, the time warps ψ are intrinsically diffeomorphic.

Remark 2 (On the noise model) In (5), we assumed the noise of the data to be Gaussian. This choice leads to the squared distance between the deformed template and each observation in the criterion (6). The same assumption is made by the registration schemes that are driven by "sum of squared differences"-like metrics. Though convenient, this noise model is arguable. In the framework of currents, the simulation of a Gaussian noise is equivalent to adding random Dirac delta currents at the nodes of a regular lattice, whose covariance matrix is given by the kernel (momenta close to each others tend to be correlated), which is not unlike a sensor noise (see Durrleman 2010, Chap. 3 for more details). Other noise models that are more closely related to the mesh

structure of the surfaces could be used, at the cost of a more complex MAP derivation.

Note that, in our model, we supposed the subjects' data to be corrupted by noise, and not the template which is supposed to be a noise free ideal representation of the shape. Therefore, the resulting cost function (6) is not symmetric, as the observed subjects' shapes do not play the same role as the template.

3.2 Computational Framework and Algorithms

3.2.1 A Generic Way to Build Diffeomorphisms

In this section, we explain a way to build generic 3D and 1D flows of diffeomorphisms which will be used as a model for the deformations χ_t, ϕ and ψ in the following. We use here the LDDMM framework (Trouvé 1998; Dupuis et al. 1998; Miller et al. 2002) for constructing 3D diffeomorphisms. We propose to adapt this framework to the construction of $1D$ diffeomorphisms.

3D diffeomorphisms In the LDDMM framework, 3D diffeomorphisms are generated by integrating time-varying vector fields. Let $v_t(x)$ be a time-varying speed vector field which gives the velocity of a particle which is at position x at time t. A particle which is at position x at time $t = 0$ moves to the position $\phi_t(x)$ at time t. The function $\phi_t(x)$ follows the differential equation for $t \geq 0$:

$$\begin{cases} \frac{d\phi_t(x)}{dt} = v_t(\phi_t(x)) \\ \phi_0(x) = x \end{cases} \tag{7}$$

Under some conditions on the regularity of the speed vector field explained in Trouvé (1998), the set of deformation ϕ_t is a flow of diffeomorphisms of the 3D domain. Following this theory, we assume that the speed vector field belongs to a reproducible kernel Hilbert space (RKHS), meaning the speed vector fields result from the convolution between a square integrable vector field and a smoothing kernel K, which plays the role of a low-pass filter. In our applications, we will use a Gaussian kernel, which writes $K(x, y) = \sigma^2 \exp\left(-|x - y|^2 / \lambda^2\right) \mathrm{I}$ for any points (x, y) in space and I the identity matrix. The spatial scale λ determines the typical scale at which points in space have a correlated speed, and therefore move in a consistent way. It determines the degree of smoothness of the deformations. Large scale means almost rigid deformations. Small scales favor deformations with many small-scale local variations. The parameter σ is a scaling factor, which in some cases cancels out with the trade-offs in the criterion, as we will discuss later.

In this setting, we define the measure of regularity of the flow of diffeomorphisms as the total kinetic energy of the flow between $t = 0$ and $t = T$:

$$\mathrm{Reg}(\phi) = \int_0^T \|v_t\|_V^2 \, dt \tag{8}$$

where $\|.\|_V$ denotes the RKHS norm associated to the kernel K.

An important property, (proven in Joshi and Miller (2000),Glaunès (2005) and extended in Durrleman (2010), Chap. 4 in case of the matching term involves several time-points), states that the vector field in the RKHS V which achieves the best trade-off between this regularity term and a fidelity-to-data term has a finite-dimensional parameterization, if the fidelity-to-data term depends only on a finite number of points:

Proposition 1 (Finite dimensional parameterization of minimizing vector field) *Let E be a criterion of the form:*

$$E(v) = \sum_i A_i(\phi_{t_i}^v(S)) + \gamma \int_0^T \|v_t\|_V^2 \, dt \tag{9}$$

where v_t denotes a time-varying speed vector field, ϕ_t^v the flow generated by this vector field in the sense of (7), S a discrete set of N points x_i in the 3D domain and A_i a set of positive and continuous functions from $\mathbb{R}^{3N}$ to $\mathbb{R}$.

Then, the criterion E admits at least one minimum and the vector field which minimizes E over all possible vector field in the RKHS V is parameterized by a set of N time-varying vectors $(\alpha_i(t))$, such that:

$$v_t(x) = \sum_{i=1}^N K(x, x_i(t))\alpha_i(t) \tag{10}$$

for any points x, where $x_i(t) = \phi_t^v(x_i)$ satisfies the flow equations:

$$\begin{aligned} \frac{dx_i(t)}{dt} &= v_t(x_i(t)) \\ &= \sum_{j=1}^N K(x_i(t), x_j(t))\alpha_j(t) \quad \text{with} \quad x_i(0) = x_i \end{aligned} \tag{11}$$

The couples $(x_i(t), \alpha_i(t))$ are called momenta.

The norm of the minimizing vector field in the RKHS V is given as:

$$\|v_t\|_V^2 = \sum_{i=1}^N \sum_{j=1}^N \alpha_i(t)^t K(x_i(t), x_j(t))\alpha_i(t), \tag{12}$$

The criterion depends therefore only on the set of L^2 functions $\alpha_i(t)$. Given these functions and the initial positions x_i, one can integrate (11) to generate the trajectories $x_i(t)$. Then the criterion E involves only the couples $(x_i(t), \alpha_i(t))$.

1D diffeomorphisms We can adapt this framework to the construction of $1D$ diffeomorphisms, which will be used as

time warps in our method. Let the variable $t \in \mathbb{R}$ play the role of the spatial variable x in the construction of 3D diffeomorphisms. We can build a flow of 1D diffeomorphisms $\psi_u(t)$ for the parameter u in [0, 1] (here u plays the previous role of t, since now t denotes a 'real' time and not the integration variable) by integrating the flow equation:

$$\begin{cases} \frac{d\psi_u(t)}{du} = v_u(\psi_u(t)) \\ \psi_0(t) = t \end{cases} \tag{13}$$

where v_u is now a scalar function, which gives the speed at which the time $\psi_u(t)$ evolves. If it is positive, time tends to accelerate. If it is negative, time tends to slow down. We impose that v_u is in 1D RKHS, determined by the kernel $K(t,t') = \sigma^2 \exp\left(-\left|t-t'\right|^2/\lambda^2\right)$. The scalar parameter λ determine the typical time-length at which two time-points t and t' are changed in a correlated manner. An illustration of the construction of such 1D diffeomorphism is given in Fig. 10.

Then, the same property as Proposition 1 applies. If E is a criterion of the form:

$$E(v) = A(\psi_1(\mathbf{t})) + \gamma \int_0^1 \|v_u\|_V^2 \, du \tag{14}$$

where $\mathbf{t}$ denotes a vector of time-points $t_1, \ldots, t_N$ and A a positive and continuous scalar function, then the minimum of E over the RKHS exists and is achieved for a speed function v_u of the form:

$$v_u(t) = \sum_{i=1}^{N} K(t, \psi_u(t_i))\beta_i(u) \tag{15}$$

where the time-varying scalars $\beta_i(u)$ are L^2 functions from [0, 1] to $\mathbb{R}$. The norm of the speed function v_u in the RKHS is given by:

$$\|v_u\|_V^2 = \sum_{i=1}^{N}\sum_{j=1}^{N} K(\psi_u(t_i), \psi_u(t_j))\beta_i(u)\beta_j(u) \tag{16}$$

3.2.2 Optimization of the Regression Criterion

We optimize the regression criterion (1) assuming that the regression function χ_t is generated by a time-varying velocity field v^χ, which belongs to the RKHS V^χ determined by the 3D Gaussian kernel K^χ with standard deviation λ_χ. Defining the regularity criterion as $\mathrm{Reg}(\chi) = \int_0^T \|v_t\|_{V^\chi}^2 \, dt$, the criterion to be minimized becomes:

$$E(\chi) = \sum_{t_i} \mathrm{d}(\chi_{t_i}(S_0), S_i)^2 + \gamma^\chi \int_0^1 \|v_t\|_{V^\chi}^2 \, dt \tag{17}$$

where d is a similarity measure between shapes. Here we assume that the baseline S_0 and the shape S_i are sets of points, polygonal lines or meshes. We denote $(x_1, \ldots, x_N)$ the vertices of the baseline S_0. Would d be either the sum of squared differences between point positions or the distance on currents in absence of point correspondence, the conditions of Proposition 1 are satisfied (Glaunès 2005): the minimizing vector field v_t is parameterized by the momenta $(x_p(t), \alpha_p(t))$: $v_t^\chi(x) = \sum_{p=1}^N K^\chi(x, x_p(t))\alpha_p(t)$.

As noticed in Sect. 3.2.1, the regression criterion E is a function of the N L^2 functions $\alpha_p(t)$. We provide this set of N functions with the metric induced by the kernel K^χ, meaning that the inner-product between two sets of L^2 functions $\alpha_p(t)$ and $\alpha'_p(t)$ is given by:

$$\int_0^1 \sum_{p=1}^{N}\sum_{q=1}^{N} (\alpha'_p(t))^t K^\chi(x_p(t), x_q(t))\alpha_q(t)dt.$$

The gradient of E with respect to the p*th* function $\alpha_p(t)$ is an L^2 function denoted $\nabla_{\alpha_p} E(t)$, which is such that for all:

$$\frac{d}{d\tau} E(\alpha_1(t), \ldots, \alpha_p(t) + \tau\varepsilon_p(t), \ldots, \alpha_N(t)) = \int_0^1 \sum_{q=1}^{N} \varepsilon_p(t)^t K^\chi(x_p(t), x_q(t))\nabla_{\alpha_q} E(t)dt.$$

In "Appendix A", we show that this gradient is equal to:

$$\nabla_{\alpha_p} E(t) = 2\gamma^\chi \alpha_p(t) + \eta_p^\chi(t) \tag{18}$$

where $\eta_p^\chi(t)$ is the solution of the linear set of backward integral equations for all p:

$$\begin{aligned} \eta_p(t) = &\sum_{t_i} (\nabla_{x_p(t_i)} A_i)\mathbf{1}_{\{t \le t_i\}} \\ &+ \int_t^T \sum_{q=1}^{N} \Big(\alpha_p(u)^t \eta_q(u) + \alpha_q(u)^t \eta_p(u) \\ &+ 2\gamma^\chi \alpha_p(u)^t \alpha_q(u)\Big)\nabla_1 k^\chi(x_p(u), x_q(u))du \end{aligned} \tag{19}$$

where $\mathbf{1}_{\{t \le t_i\}} = 1$ if $t \le t_i$ and 0 otherwise, $A_i = \mathrm{d}(\chi_{t_i}(S_0), S_i)^2$ seen as a function of the points positions $x_p(t_i)$ and $k^\chi(x,y) = \exp\left(-|x-y|^2/\lambda_\chi^2\right)$.

The gradient descent scheme for the computation of the regression $S(t) = \chi_t(S_0)$ is summarized in Algorithm 1 in "Appendix C". We start the gradient descent by setting $\alpha_p(t) = 0$ for all t and p ($\chi_t = \mathrm{Id}$ and $S(t) = S_0$, for all t). Computing the gradient requires first to integrate of the flow equation (Eq. 11) forward in time and then to compute the auxiliary variable η^χ (Eq. 19) backward in time. In this last case, the initial conditions at $t = T$ is given by $\nabla_{x_p(T)} A_T$. Then the ODE is integrated for decreasing time t. As soon as a new time point t_i is reached, a new contribution $\nabla_{x_p(t_i)} A_i$ is added to $\eta^\chi(t)$. As a consequence, $\nabla_{\alpha_p} E(t)$ (and therefore

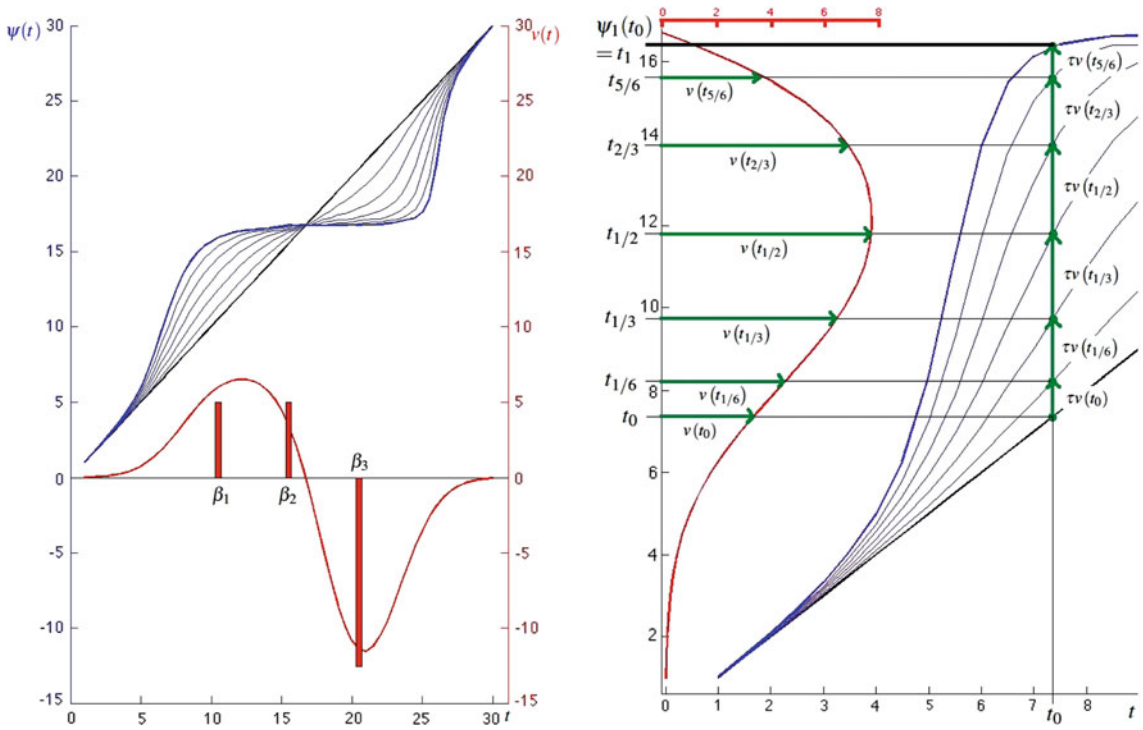


Fig. 10 Construction of 1D diffeomorphisms by integration of speed functions. In this illustration, we suppose the speed function to be constant (v independent of u): $\frac{\partial \psi_u(t)}{dt} = v(\psi_u(t))$. *Left* The speed profile v is set as the convolution of 3 constant momenta (β_i) with a Gaussian kernel with standard deviation $\lambda_\psi = 4$ (in *red*). The integration of the flow equation with the initial condition $\psi_0(t) = t$ is shown in *blue*: the *bold blue curve* corresponds to the final diffeomorphism at $u = 1$, *light blue curves* correspond to $\psi_{1/6}(t)$, $\psi_{1/3}(t)$, $\psi_{1/2}(t)$, $\psi_{2/3}(t)$ and $\psi_{5/6}(t)$. *Right* Illustration of the numerical integration of the flow: $\psi_{u_{n+1}}(t) = \psi_{u_n}(t) + \tau v\left(\psi_{u_n}(t)\right)$. The speed profile in red is shown along the y-axis. One can show easily that this scheme produces only increasing function (invertible 1D function), when τ is chosen small enough

the momenta $\alpha_p(t)$ and the vector field v_t^χ) at time t depend on all the data which appear later than t. Once the momenta are updated, the new positions $x_p(t)$ are computed by the integration of the flow Eq. (11) forward in time (the initial condition is given at time $t = 0$ by $x_p(0) = x_p$). These positions at time t depend on the vector field v_t^χ for all time earlier than t. As a result, the positions $x_p(t)$ depend on all the data in past and future. This regression fits the best trajectory ($\chi_t(S_0)$) to all the data *globally*. This differs, for instance, from pairwise registrations between consecutive time-points, although both techniques result in a piecewise geodesic flow.

For better numerical accuracy, we replace the Euler scheme in Algorithm 1 to integrate ODEs by a Euler scheme with prediction/correction, which has the same accuracy as a Runge–Kutta method of order 2. The computational bottleneck of this algorithm is the computation of every sum of the form $\sum_{p=1}^N K(x_q, x_p)\alpha_p$ that need to be computed for all q. These computations of complexity N^2 (where N is number of points in the baseline shape) can be efficiently approximated using a linearly spaced grid and FFT (Durrleman 2010), or Fast Multipole Approximations (Glaunès 2005), with a nearly linear complexity.

The computation of the gradient requires to compute the differentiation of the fidelity-to-data term: $\nabla_{x_p(t_i)} d(\phi_{t_i}(S_0) - S_i)^2$. If S_i have the same number of points as S_0 (i.e. N points), then d can be defined as the sum of squared differences: $A_i = \sum_{p=1}^N \left| x_p(t_i) - s_p^i \right|^2$, where s_p^i denotes the points of S_i. In this case, $\nabla_{x_p(t_i)} A_i = 2(x_p(t_i) - s_p^i)$. In absence of point correspondence, the distance on currents is used, which can be differentiated as explained in Glaunès (2005) and Durrleman (2010).

3.2.3 Optimization of the Spatiotemporal Registration Criterion

As explained in Sect. 3.2.1, the morphological deformation ϕ and the time warp ψ are generated by the integration of flows of 3D and 1D velocity fields respectively. This means that they are the end-points $\phi = \phi_1$ and $\psi = \psi_1$ of the differential equations:

$$\begin{aligned} \frac{\partial \phi_u(x)}{\partial u} &= v^\phi(\phi_u(x)) \\ \frac{\partial \psi_u(t)}{\partial u} &= v^\psi(\psi_u(t)) \end{aligned} \tag{20}$$

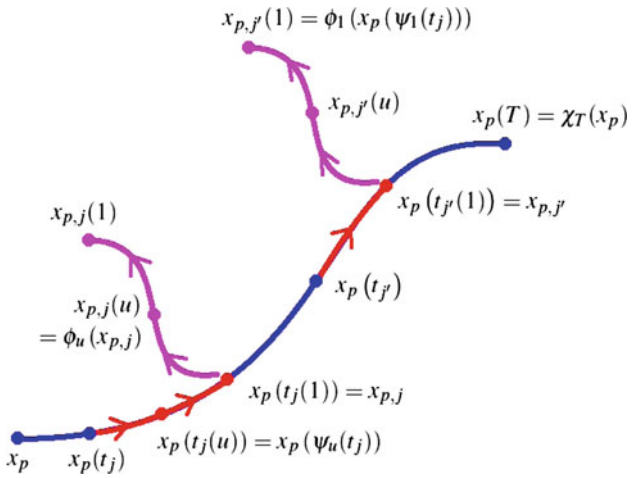


Fig. 11 Illustrative scheme for the notations: x_p denotes a generic point of the source shape, $x_p(t) = \chi_t(x)$ the continuous evolution of the source point, $(\psi_u)_{u\in[0,1]}$ is a flow of $1D$-diffeomorphism which moves the time-labels along the time-axis (in *red*), $(\phi_u)_{u\in[0,1]}$ is a flow of $3D$-diffeomorphism which moves the points of the source evolution (in *magenta*), *independently* at each time-point

with the initial conditions: $\phi_0(x) = x$ and $\psi_0(t) = t$.

Now, we assume that for every parameter u, the 3D velocity fields v_u^ϕ and 1D velocity profile v_u^ψ belong to a RKHS with Gaussian kernel K^ϕ and K^ψ, with standard deviation λ_ϕ and λ_ψ respectively. We denote $\{x_p\}_{p=1,\dots,N}$ the set of points of the discrete shape S_0. The source trajectory $S(t) = \chi_t(S_0)$ is described by the moving points $x_p(t)$. Let $\{t_j\}_{j=1,\dots,N_{\text{target}}}$ be the time-points associated to target shapes.

The fidelity-to-data term in (4) depends on the variables $\phi_1\left(x_p(\psi_1(t_j))\right) = \phi_1(x_{p,j})$, where we denote (see Fig. 11 for an illustrative scheme):

$$x_{p,j} = x_p(\psi_1(t_j)) \tag{21}$$

Therefore, the application of Proposition 1 leads to the following parameterization of the minimizing velocity fields:

$$v_u^\phi(x) = \sum_{p=1}^{N}\sum_{j=1}^{N_{\text{target}}} K^\phi(x, \phi_u(x_{p,j}))\alpha_{p,j}(u) \tag{22}$$

and

$$v_u^\psi(t) = \sum_{j=1}^{N_{\text{target}}} K^\psi(t, \psi_u(t_j))\beta_j(u) \tag{23}$$

The criterion (4) is a function of the $N \times N_{\text{target}}$ L^2 functions $\alpha_{p,j}$ and the N_{target} L^2 functions β_j. Like for the regression case, we provide this set of functions with the metric induced by the kernel K^ϕ and K^ψ.

For the sake of simplicity, we introduce the notations $x_{p,t}(u)$ and $t_j(u)$ such that:

$$\begin{aligned}
&\phi_u(x_{p,j}) = x_{p,j}(u)\\
&x_{p,j}(0) = x_{p,j} = x_p(\psi_1(t_j))\\
&x_{p,j}(1) = \phi(x_{p,j})\\
&\psi_u(t_j) = t_j(u)\\
&t_j(0) = t_j\\
&t_j(1) = \psi(t_j)
\end{aligned} \tag{24}$$

The regularity parameters in (4) are given by:

$$\begin{aligned}
\text{Reg}(\phi) &= \int_0^1 \left\|v_u^\phi\right\|_{V^\phi}^2 du\\
&= \int_0^1 \sum_{p,j,p',j'} \alpha_{p,j}(u)^t K^\phi(x_{p,j}(u),\\
&\qquad x_{p',j'}(u))\alpha_{p',j'}(u)du
\end{aligned} \tag{25}$$

and

$$\begin{aligned}
\text{Reg}(\psi) &= \int_0^1 \left\|v_u^\psi\right\|_{V^\psi}^2 du\\
&= \int_0^1 \sum_{j,j'} \beta_j(u)^t K^\psi(t_j(u), t_{j'}(u))\beta_{j'}(u)du
\end{aligned} \tag{26}$$

As shown in "Appendix B", the gradient of the criterion with respect to the functions $\alpha_{p,j}(u)$ (denoted $\nabla_{\alpha_{p,j}}E(u)$) and to the functions $\beta_j(u)$ (denoted $\nabla_{\beta_j}E(u)$) is given by:

$$\begin{aligned}
&\nabla_{\alpha_{p,i}}E(u) = 2\gamma^\phi\alpha_{p,i}(u) + \eta_{p,i}(u)\\
&\nabla_{\beta_j}E(u) = 2\gamma^\psi\beta_j(u) + \xi_j(u)
\end{aligned} \tag{27}$$

where $\eta_{i,p}(u)$ satisfies the backward integral equation:

$$\begin{aligned}
\eta_{p,i}(u) = \nabla_{x_{p,i}(1)}A + \int_u^1 \sum_{q=1}^{N}\sum_{j=1}^{N_{\text{target}}}\\
\times\Big(\alpha_{p,i}(s)^t\eta_{q,j}(s) + \eta_{p,i}(s)^t\alpha_{q,j}(s)\\
+2\gamma^\phi\alpha_{p,i}(s)^t\alpha_{q,j}(s)\Big)\nabla_1 k^\phi(x_{p,i}(s), x_{q,j}(s))ds
\end{aligned} \tag{28}$$

where A denotes the matching term $\sum_{j=1}^{N_{\text{target}}} d(\phi_1(S(\psi_1(t_j))) - U_j)^2$ which is a function of the variables $\phi_1(x_p(\psi_1(t_j))) = x_{p,j}(1)$.

and where $\xi_j(u)$ satisfies the backward integral equation:

$$\xi_j(u) = \sum_{p=1}^{N} \left(\frac{dx_p(t)}{dt}\Big|_{t=t_j(1)} \right)^t \eta_{p,j}(0) + \int_u^1 \sum_{k=1}^{N_{\text{target}}} \Big(\beta_j(s)^t \xi_k(s) + \xi_j(s)^t \beta_k(s) + 2\gamma^\psi \beta_j(s)^t \beta_k(s) \Big) \nabla_1 k^\psi(t_j(s), t_k(s)) ds \quad (29)$$

The auxiliary space variable $\eta(u)$ pulls the gradient of the matching term from $u = 1$ back to $u = 0$ along the space axis. Then, the value $\eta(0)$ is used in the final conditions of the auxiliary time variable $\xi(u)$ in combination with the local speed of the source growth scenario, thus showing the spatiotemporal coupling. The variable $\xi(u)$ pulls back this condition at $u = 1$ back to $u = 0$ along the time axis. The gradient transports the driving force in the target space, namely the gradient of the data term, back to source space along the spatiotemporal deformation. This transport is used to update the momenta $\alpha_{p,j}(u)$ and $\beta_j(u)$, which parameterize the spatiotemporal deformation.

In (28), the gradient of the matching term is computed as for the regression function. For instance, if the distance between source and target is the sum of squared differences: $A = \sum_{j=1}^{N_{\text{target}}} \sum_{p=1}^{N} |x_{p,j}(1) - U_{p,j}|^2$, then the gradient is simply $\nabla_{x_{p,i}(1)} A = 2(x_{p,j}(1) - U_{p,j})$.

In (29), one needs to compute the speed of the source growth scenario: $\frac{dx_p(t)}{dt}$. If one has stored the parameterization of the regression function (i.e. the momenta $(x_p(t), \alpha_p^\chi(t))$), then one can compute explicitly:

$$\frac{dx_p(t)}{dt} = \frac{d\chi_t(x_p)}{dt} = v_t^\chi(x_p(t)) = \sum_{q=1}^{N} K^\chi(x_p(t), x_q(t)) \alpha_q^\chi(t) \quad (30)$$

In our implementation, we only stored samples of the trajectories $x_p(t)$ and not the vectors $\alpha^\chi(t)$. So, we estimate this speed by a finite difference scheme: $\frac{dx_p(t)}{dt}(t_j) = (x_p(t_{j+1}) - x_p(t_{j-1}))/2$. This allows us to still use this spatiotemporal registration even if the source evolution has been computed with another regression method than the one presented in Sect. 3.2.2.

The sketch of the gradient descent for this spatiotemporal registration scheme is given in Algorithm 2 in "Appendix C". Note that we minimize the criterion with respect to the geometrical and the temporal parameters *jointly*, thus avoiding alternated minimization. The differentiation of the criterion proposed here is here, whereas an approximation was involved in Durrleman et al. (2009b) and Durrleman (2010). In practice, both differentiations leads to similar results.

Remark 3 Note that since the growth model χ_t is piecewise geodesic, the evolution $S(t)$ generated by χ_t is not differentiable at the time-points t_j: the continuous $S(t)$ may have different left and right derivatives. This point is discussed in depth in Durrleman (2010), where an alternative optimization procedure is proposed, which ensures that an extremum of the registration criterion is achieved at convergence, even in presence of discontinuous velocities. Another way to address the problem is to use the twice-differentiable growth model proposed in Fishbaugh et al. (2011).

3.2.4 Optimization of the Criterion for Atlas Construction

The estimation of the 4D-atlas relies on one regression function χ_t and N_{subj} spatiotemporal deformations (ϕ^s, ψ^s) for $s = 1, \ldots, N_{\text{subj}}$, where N_{subj} is the number of subjects. We use the framework of Sect. 3.2.1 to construct the 3D diffeomorphisms χ and ϕ^s and the 1D diffeomorphisms ψ^s. As a consequence, every deformation satisfies a flow equation as follows:

$$\begin{aligned} \frac{\partial \chi_t(x)}{\partial t} &= v_t^\chi(\chi_t(x)), \qquad t \in [0, T] \\ \frac{\partial \phi_u^s(x)}{\partial u} &= v_u^{\phi^s}(\phi^s(x)), \qquad u \in [0,1], s = 1, \ldots, N_{\text{subj}} \\ \frac{\partial \psi_u^s(t)}{\partial u} &= v_u^{\psi^s}(\psi^s(t)), \qquad u \in [0,1], s = 1, \ldots, N_{\text{subj}} \end{aligned} \quad (31)$$

where we suppose that the velocity fields v^χ (resp. v^{ϕ^s} and v^{ψ^s}) belong to a RKHS V^χ (resp. V^ϕ and V^ψ) determined by the Gaussian kernel K^χ (resp. K^ϕ and K^ψ) with standard deviation λ_χ (resp. λ_ϕ and λ_ψ).

We suppose that the prototype shape M_0 (to be estimated) is given by a finite set of points $\{x_p\}$. In this case, the application of Proposition 1 leads to the parameterization of the time-varying velocity fields by momenta as follows:

$$\begin{aligned} v_t^\chi(x) &= \sum_p K^\chi\left(x, x_p(t)\right) \alpha_p^\chi(t) \\ v_u^{\phi^s}(x) &= \sum_{p,t_j^s} K^\phi\left(x, x_{p,j}^s(u)\right) \alpha_{p,j}^s(u), \text{ for } s = 1, \ldots, N_{\text{subj}} \\ v_u^{\psi^s}(t) &= \sum_{t_j^s} K^\psi\left(t, t_j^s(u)\right) \beta_j^s(u), \text{ for } s = 1, \ldots, N_{\text{subj}} \end{aligned} \quad (32)$$

where we denote:

$$x_p(t) = \chi_t(x_p) \quad x_{p,j}^s(u) = \phi_u^s(x_p^s(t_j^s(1))) \quad t_j^s(u) = \psi_u^s(t_j^s) \quad (33)$$

for all $t \in [0, T]$ and $u \in [0, 1]$. t_j^s denotes the N_{target}^s time-points at which the s*th* subject has been observed, which might be different for every subject.

The criterion for atlas estimation depends therefore on the N points of the template $M_0 = \{x_p\}_{p=1,\dots,N}$, the N_{time} t-varying vectors $\alpha_p(t)$ for the regression function, the $N \times \sum_{s=1}^{N_{\text{subj}}} N^s_{\text{target}}$ u-varying vectors $\alpha^s_{p,j}(u)$ for the morphological deformations and the $\sum_{s=1}^{N_{\text{subj}}} N^s_{\text{target}}$ u-varying vectors $\beta^s_j(u)$ for the time warps. This criterion can be written now as:

$$J\left(\{\alpha^s_{p,j}(u)\}, \{\beta^s_j(u)\}, \{\alpha^\chi_p(t)\}, M_0\right) = \sum_{s=1}^{N_{\text{subj}}} \left\{ \sum_{s=1}^{N_{\text{subj}}} \sum_{j=1}^{N^s_{\text{target}}} A\left(\{x^s_{p,j}(1)\}\right) + \gamma^\phi \int_0^1 \left\| v^{\phi^s}_u \right\|^2_{V^\phi} du + \gamma^\psi \int_0^1 \left\| v^{\psi^s}_u \right\|^2_{V^\psi} du + \gamma^\chi \int_0^T \left\| v^\chi_t \right\|^2_{V^\chi} dt \right\} \tag{34}$$

where the matching term $A\left(\{x^s_{p,j}(1)\}\right) = \mathrm{d}(\phi^s(\chi_{\psi^s(t^s_j)} M_0), S^s_j)^2$ depends on the positions $\phi^s_1(M(\psi_1(t^s_j))) = \{(x^s_{p,j}(1)\}_{p=1,\dots,N}$.

To minimize this criterion, we adopt a 3-step alternating minimization procedure:

- If the template M_0 and the growth function χ_t are fixed, the criterion is divided into N_{subj} independent functions. Their minimum is achieved for the spatiotemporal deformations (ϕ^s, ψ^s), which maps the mean scenario $\chi_t(M_0)$ to the set of data S^s_j for each subject s. These N_{subj} spatiotemporal registrations are computed using Algorithm 2.
- If the N_{subj} spatiotemporal deformations (ϕ^s, ψ^s) and the growth function χ_t are fixed, the criterion to be minimized with respect to M_0 is reduced to:

$$J(M_0) = \sum_{s,j} \mathrm{d}(\Phi_{s,j}(M_0), S^s_j)^2,$$

where we denote $\Phi_{s,j} = \phi^s \circ \chi_{\psi^s(t^s_j)}$. These deformations are $3D$-diffeomorphisms. This criterion has exactly the form of the criterion for usual 3D template estimation. If d is the distance on currents, a solution for the minimization of this convex criterion has been proposed in Durrleman et al. (2009a) and Durrleman (2010, Chap. 5, Algorithm 4). As a consequence, the template M_0 is always given as a finite set of points $\{x_p\}_{p=1,\dots,N}$.
- If the template M_0 and the N_{subj} spatiotemporal deformations (ϕ^s, ψ^s) are fixed, the criterion to be minimized becomes:

$$\sum_{s,j} \mathrm{d}(\phi^s(\chi_{\psi^s(t^s_j)} M_0), S^s(t^s_j))^2 + \gamma^\chi \mathrm{Reg}(\chi).$$

This is not exactly the regression problem stated in Sect. 3.1.1 because of the deformation ϕ^s in the matching term. To turn it into a regression problem, we approximate the matching term $\mathrm{d}(\phi^s(\chi_{\psi^s(t^s_j)} M_0), S^s(t^s_j))$ by $\mathrm{d}(\chi_{\psi^s(t^s_j)}(M_0), (\phi^s)^{-1}(S^s(t^s_j)))$, meaning that the shapes of each subject are matched back to the mean anatomy. This approximation is valid only for diffeomorphisms ϕ^s whose Jacobian is close to the identity, since the usual metrics d are not left-invariant. As a result, the evolution function χ_t performs the temporal regression of the set of shapes $(\phi^s)^{-1}(S^s_j)$ located at time-points $\psi^s(t^s_j)$. This regression problem can now be solved using Algorithm 1. Further investigations are needed in order to perform this regression without this approximation, so that we can be consistent throughout the minimization procedure.

To initialize the minimization, we set M_0 as the mean current of the earliest data ((S^s_1) for every subject s) and set all the momenta to zeros (χ, ϕ^s, ψ^s equal identity map). The whole minimization procedure is summarized in Algorithm 3 in "Appendix C".

3.2.5 Parameters

The overall framework depends on several parameters. There are 3 kernels of 3 distinct RKHS: K^χ, K^ϕ and K^ψ. We use Gaussian kernels determined by their standard deviations: λ_χ, λ_ϕ and λ_ψ respectively. They determine the degree of smoothness (i.e. the scale at which points have a correlated speed) of the mean scenario of evolution, the morphological deformations and the time warp. The first one compares with the scale of the geometrical variations of the structure over time for a typical subject (scale of the intra-subject variability). The second one compares with the scale of geometrical variations between different subjects (geometrical inter-subject variability). The third one compares with the typical time-scale at which the dynamics of evolution changes from subject to subject.

The user must also set the 3 trade-offs between regularity and fidelity to data: γ^χ, γ^ϕ, γ^ψ. In addition, one needs to set the metric d between shapes. In the framework of currents, this metric depends on a kernel K^W. We choose a Gaussian kernel with standard deviation λ_W. This parameter sets the typical scale at which shape variations are smoothed (see Durrleman (2010)).

The dimension of the trade-off γ^χ, γ^ϕ and γ^ψ depends on the kinds of data that we deal with. The dimension of the data term in the criterions is L^2 (i.e. squared length) for curves and L^4 (i.e. squared area) for surfaces, where L denotes the dimension of a length. The parameter t has the dimension of time (denoted T) and the parameter u is an integration

parameter, which is normalized to fall in the unit interval [0, 1] and therefore has no physical dimension. Therefore, the velocities v_t^χ, v_u^ϕ, v_u^ψ are of dimension LT^{-1}, L and T respectively and the regularity terms (integral of the squared norm of the velocities) $\mathrm{Reg}(\chi)$, $\mathrm{Reg}(\phi)$ and $\mathrm{Reg}(\psi)$ are of dimension: L^2T^{-1}, L^2 and T^2 respectively. Eventually, the dimension of the trade-off γ is that of the ratio between the data term and the regularity terms:

	Curves	Surfaces
γ_χ	T	L^2T
γ_ϕ	None	L^2
γ_ψ	L^2T^{-2}	L^4T^{-2}

In the future, we plan to normalize these constants so that their values can be compared for different applications. More generally, one needs to better understand the balance between the spatial and temporal constraints and to find an automatic way to estimate this parameters (which could be considered as fixed effects in a Bayesian framework along the lines of Allassonnière and Kuhn (2009)).

3.3 Statistical Measures of Spatiotemporal Variability

The construction of the spatiotemporal atlas leads to the mean scenario $M(t)$, which gives a representative mean of the studied population, and the spatiotemporal deformations of this mean scenario to each subject, which estimates the variance within the population. The criterion for the atlas construction is not unlike the estimation of a Fréchet mean on the "manifold" of the individual trajectories, the distance between two individual trajectories being given by the cost of the spatiotemporal deformation which connects them (see Miller et al. (2002) for this interpretation in the 3D case). In this section, we explain how one can compute *intrinsic* statistics on the spatiotemporal deformations: the mean and the principal modes of the morphological deformations and the time warps are defined as 3D and 1D diffeomorphisms themselves. Due to the definition of the mean scenario, the mean of all deformations vanishes. Nevertheless, one can compute the means of population sub-groups to detect significant differences between them. The modes show the typical variations of the mean scenario within the population or within one sub-group. They can be used to drive the search for anatomical characterization of sub-groups. Besides the quantification of group differences and the usual hypothesis testing, one important aspect of intrinsic statistics is that means and modes can be displayed as movies of shape evolutions, which is crucial for qualitative interpretation purposes. This can be used to better understand the effect of a pathology and drive the search for bio-markers.

3.3.1 Statistics on Initial Momenta

As shown in Miller et al. (2002), the flow of diffeomorphisms which minimize the registration criterion (4) or the atlas construction criterion (6) are geodesic: they are the ones which minimize the length of the path $(\phi_u, \psi_u)_{u\in[0,1]}$ between the identity map Id and the actual diffeomorphisms (ϕ_1, ψ_1). These initial velocity plays the role of a tangent-space representation as in finite-dimensional Riemannian geometry (Pennec et al. 2006): they are the equivalent of the logarithm of the deformations. Since we perform template-to-subjects registration (and not subjects-to-template), every flows of diffeomorphisms ϕ^s (resp. ψ^s) starts from the same space, the one of the mean scenario, and therefore share the same tangent-space V^ϕ (resp. V^ψ). As a consequence, one can perform intrinsic statistics on these common vector spaces. Since the initial velocities are parameterized by a finite number of momenta, the statistics on deformations reduces to statistics in an Euclidean space.

For each subject s, the 3D diffeomorphisms ϕ^s are parameterized by momenta $\alpha^s_{p,j}(0)$ located at the points $x_{p,j}(0) = x_p(t_j^s(1))$, which is a subset of the whole point set $x_{p,k}$ (the trajectories of every template point). Using zero-padding, every ϕ^s is parameterized by a vector of the same dimension $\boldsymbol{\alpha}^s = \{\alpha^s_{p,k}\}_{p,k}$. Similarly, each 1D diffeomorphism ψ^s is characterized by momenta $\beta_j^s(0)$ located at time-points $t_j^s(0) = t_j^s$, which is a subset of the set of all time-points $\{t_k\}$. Using zero-padding, every 1D diffeomorphism is characterized by a vector of the same dimension: $\boldsymbol{\beta}^s = \{\beta_k^s\}_k$.

One can compute a Principal Component Analysis (PCA) on the vectors $\boldsymbol{\alpha}^s$ and $\boldsymbol{\beta}^s$ according to the metric on the RKHS V^ϕ and V^ψ as follows (see Durrleman 2010, Chap. 5 for more details). One builds the mean vectors $\overline{\boldsymbol{\alpha}} = \sum_s \boldsymbol{\alpha}^s / N_{\mathrm{subj}}$ and $\overline{\boldsymbol{\beta}} = \sum_s \boldsymbol{\beta}^s / N_{\mathrm{subj}}$ and the centered vectors $\tilde{\boldsymbol{\alpha}}^s = \boldsymbol{\alpha}^s - \overline{\boldsymbol{\alpha}}$ and $\tilde{\boldsymbol{\beta}}^s = \boldsymbol{\beta}^s - \overline{\boldsymbol{\beta}}$. Then, one builds the empirical matrices Σ^ϕ and Σ^ψ of size $N_{\mathrm{subj}} \times N_{\mathrm{subj}}$ whose term s, s' is given by:

$$\begin{aligned}
\Sigma^\phi_{s,s'} &= \left\langle \tilde{\boldsymbol{\alpha}}^s, \tilde{\boldsymbol{\alpha}}^{s'} \right\rangle_{V^\phi} \\
&= \sum_{k,k'} \sum_{p,p'=1}^{N} (\tilde{\alpha}^s_{p,k})^t K^\phi(x_{p,k}, x_{p',k'}) \tilde{\alpha}^{s'}_{p',k'} \\
\Sigma^\psi_{s,s'} &= \left\langle \tilde{\boldsymbol{\beta}}^s, \tilde{\boldsymbol{\beta}}^{s'} \right\rangle_{V^\psi} = \sum_{k,k'} K^\psi(t_k, t_{k'}) \tilde{\beta}^s_k \tilde{\beta}^s_{k'}
\end{aligned} \tag{35}$$

Let $\mathbf{E}^\phi_m$ and $\mathbf{E}^\psi_m$ the eigenvectors of the matrices Σ^ϕ and Σ^ψ associated to the m*th* largest eigenvalues λ^ϕ_m and λ^ψ_m (these are vectors of dimension N_{subj}). Then, as shown in Durrleman (2010, Chap. 5), the *m*th eigenmode is given by:

$$\begin{aligned}\alpha_m &= \overline{\boldsymbol{\alpha}} + \sum_{s=1}^{N_{\text{subj}}} \mathbf{E}^{\phi}_{m,s}(\boldsymbol{\alpha}^s - \overline{\boldsymbol{\alpha}}) \\ \beta_m &= \overline{\boldsymbol{\beta}} + \sum_{s=1}^{N_{\text{subj}}} \mathbf{E}^{\psi}_{m,s}(\boldsymbol{\beta}^s - \overline{\boldsymbol{\beta}})\end{aligned} \tag{36}$$

such that $\|\alpha_m - \overline{\alpha}\|^2_{V^\phi} = \lambda^\phi_m$ and $\|\beta_m - \overline{\beta}\|^2_{V^\psi} = \lambda^\psi_m$.

3.3.2 Geodesic Shooting for Computing Intrinsic Means and Modes

Once one has computed the statistics on the tangent-spaces V^ϕ and V^ψ, one needs to use the geodesic shooting equations (the equivalent of the exponential map in Riemannian manifold) to generate the 3D and 1D diffeomorphisms, whose initial velocities are parameterized by the computed mean or modes. The computed diffeomorphisms are called the mean or the modes of the deformations.

Given vectors $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ (located at the points $\{x_{p,k}\}_{p,k}$ and $\{t_k\}$), the diffeomorphisms, whose initial velocities are parameterized by these vectors, are re-constructed by integrating over the interval $u \in [0, 1]$:

$$\begin{cases} \frac{d\alpha_{p,k}(u)}{du} = -\left(d_{x_{p,k}(u)} v^\phi_u\right)^t \alpha_{p,k}(u) \\ \frac{dx_{p,k}(u)}{du} = v^\phi_u(x_{p,k}(u)) \end{cases} \tag{37}$$

where $v^\phi_u(x) = \sum_{p,k} K^\phi(x, x_{p,k}(u))\alpha_{p,k}(u)$ (see Miller et al. 2006 for the proof). The positions of the points $x_{p,k}(1) = \phi_1(x_{p,k})$ builds a movie, which shows the mean (if $\overline{\alpha}$ is used as initial conditions) or the m*th* mode (if $\pm\alpha_m$ is used as initial conditions) of the morphological deformations within the population or one of its sub-group.

The equivalent equations for 1D diffeomorphisms are given as:

$$\begin{cases} \frac{d\beta_k(u)}{du} = -\left(d_{t_k(u)} v^\psi_u\right)^t \beta_k(u) \\ \frac{dt_k(u)}{du} = v^\psi_u(t_k(u)) \end{cases} \tag{38}$$

where $v^\psi_u(t) = \sum_k K^\psi(t, t_k(u))\beta_k(u)$. The integration of this set of ODEs leads to the 1D diffeomorphism $t_k(1) = \psi_1(t_k)$, called the mean time warp (if the mean vector $\overline{\boldsymbol{\beta}}$ is used as initial conditions) or the m*th* mode of the time warps (if the m*th* mode $\pm\boldsymbol{\beta}_m$ is used as initial conditions).

The geodesic shooting of the mean and the principal modes of the momenta leads to diffeomorphisms. In this sense, they are "intrinsic" statistics on an infinite-dimensional "manifold" of diffeomorphisms. In particular, the mean or the modes of the time warps are all smooth monotonic functions. This differs from computing the point-by-point mean of the real-values functions $\psi^s(t)$. Examples will be shown in the next sections.

4 Measures of Developmental Delays of Deep Brain Structures in Autism

In this section, we apply the tools introduced in Sect. 3 to analyze a longitudinal database of deep brain structures segmented from images of autistic, developmental delayed and control children (Hazlett et al. 2005, 2011). Each child has been scanned twice: a baseline at about age 3 years and a follow-up at about age 5 years. The segmentation provides a set of 24 meshes for each structure: 12 subjects divided on 3 groups of 4 subjects (autistics, developmental delays and controls), each subject having two meshes (a baseline and a follow-up). As a pre-processing, all the meshes were co-registered via rigid transformations using `gmmreg` (Jian and Vemuri 2005).

Our purpose is to show how our methodology can be used to give a description of the effect of the pathology on the maturation of the hippocampus and the amygdala of the right hemisphere. Due to the limited number of subjects involved, this study is mostly a proof of concept, aiming at showing the strengths and the limitations of our approach.

4.1 Spatiotemporal Atlas Estimation

We estimate a spatiotemporal atlas by minimizing the criterion in (6) using the algorithms described in Sect. 3.2.4. We set the time-interval of interest to [0.5, 7.1] years with a time-step of 0.2 years. The parameters of the Gaussian kernels were set to $\lambda_\chi = 10\,\text{mm}$ for the regression function, $\lambda_\phi = 10\,\text{mm}$ and $\sigma_\phi = 1$ for the morphological deformation and $\lambda_\psi = 1.5\,\text{years}$ and $\sigma_\psi = 1$ for the time warp. The typical scale on currents λ_W is set to 5 mm. The trade-offs were set to $\gamma^\chi = 10^{-4}\,\text{mm}^2\text{year}$, $\gamma^\phi = 10^{-4}\,\text{mm}^2$ and $\gamma^\psi = 10^{-6}\,\text{mm}^4\text{year}^{-2}$. The diameter of the hippocampus is about 25 mm. We refer the reader to Sect. 3.2.5 for a discussion about the parameters used (an empirical study of the impact of these parameters will be presented in the next section about endocranial data).

The output of the algorithm is a prototype shape, a mean scenario of evolution of this prototype shape, and 12 spatiotemporal deformations of this mean scenario to the pair of meshes of each subject. The analysis of the value of the criterion at the minimum shows that one autistic patient has a residual significantly larger than the other subject. This subject can be considered as an outlier, as will discuss later on. As shown in Durrleman et al. (2009a), the prototype shape is given as a current, which does not correspond to a mesh. For visualization purposes, and for the following volume computations, we mapped one instance of the data to the prototype shape and used it as a template. As a consequence, the mean scenario and its spatiotemporal deformations can be seen as the continuous evolution of a mesh.

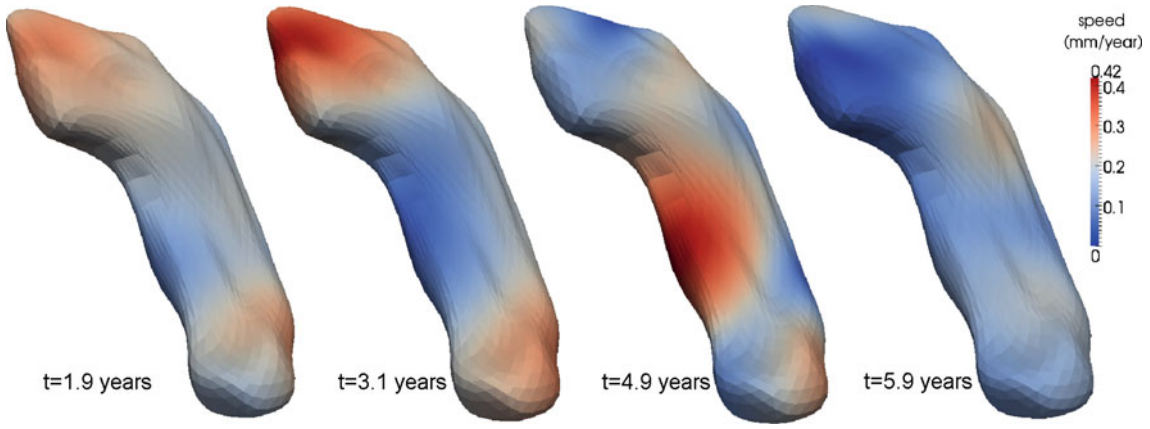


Fig. 12 Mean growth scenario of the hippocampus. Four significant frames are shown (*lateral view*). *Color* indicates the instantaneous speed of the surface deformation (best seen as a movie: see Online Resource 1)

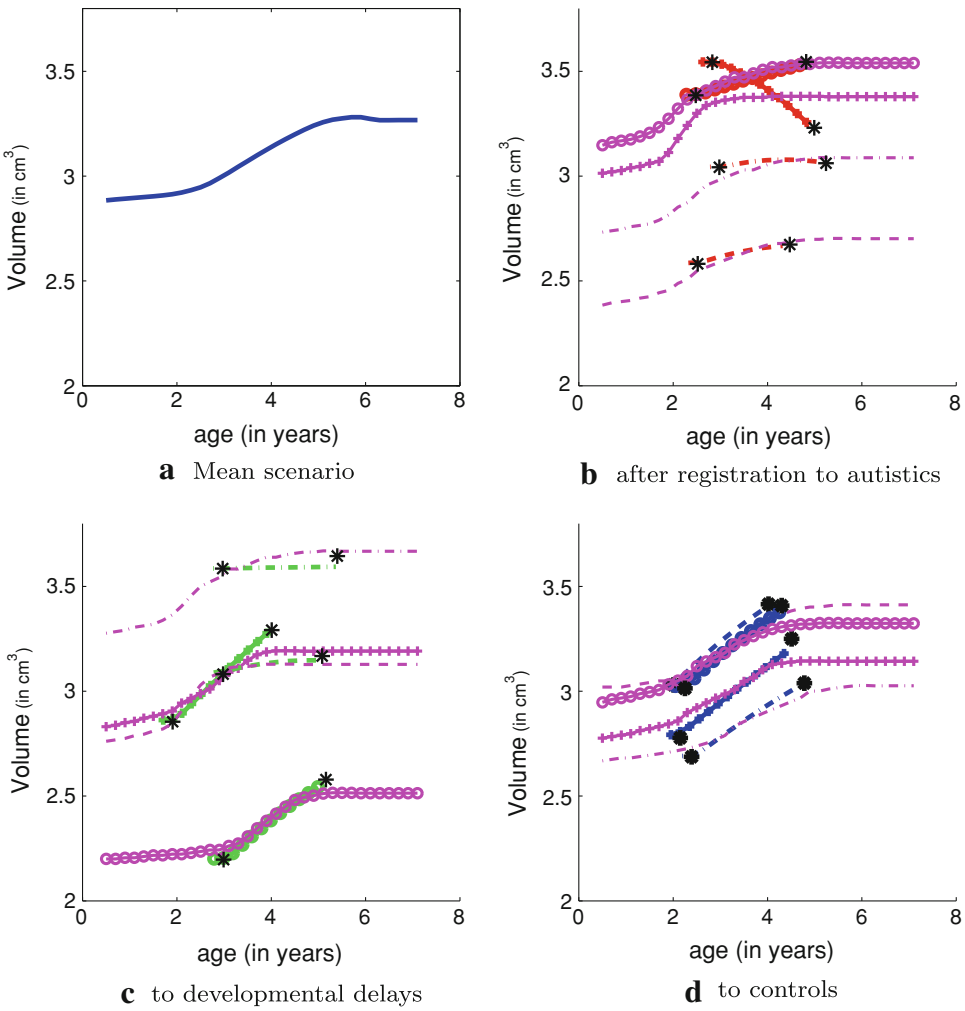


Fig. 13 Evolution of the volume of the hippocampus. (**a**) Volume evolution of the mean scenario. (**b**–**d**) Volume of the mean scenario, after it has been registered to each subject (*magenta curves*). *Black asterisks* indicate the volume of the original data. *Red*, *green* and *blue curves* indicate the volume evolution given by the pairwise registration between each subject's data pair. The autistic outlier corresponds to the decreasing *red curve* in (**b**) (decrease of volume between the two observations of this subject)

Significant samples of the estimated mean growth scenario are shown in Fig. 12. The complete scenario can be seen in the companion movie (see Online Resource 1). This mean scenario of evolution shows that the prototype growth of the structure is much more complex than a pure volume scaling over time, and involves several non-linear growth patterns. The visual inspection of the companion movie shows mainly 3 phases of growth: from 1.5 to 2.5 years, the hippocampus tends to "unfold" giving it less curved aspect; from 2.5 to 4.5 years, the hippocampus strongly elongates in the

antero/posterior direction; from 4.5 to 6 years, the extremities of the hippocampus tends to bend: the head toward the bottom and the tail toward the top, thus stretching the body of the hippocampus (red blob in the third frame in Fig. 12).

This mean scenario of evolution has been estimated along with its spatiotemporal deformation to each subject. The spatiotemporal deformation takes into account all the shape information and not only the size. Nevertheless, to illustrate the method, we compute the volume of the original meshes, of the template mesh and the deformed meshes. The evolution of the volume of the mean hippocampal growth is shown in Fig. 13a. Although the growth involves different non-linear patterns in shape as highlighted in Fig. 12, the volume extracted from this mean scenario evolves quite linearly between age 2 and 6 years. Outside this interval, the volume remains constant due to the boundary conditions (the growth function χ_t equals identity). In Fig. 13b–d, we plot the volume evolution of the mean scenario, once it has been registered to each subject, taking into account both the morphological deformation and the time warp: the morphological deformation changes the values of the curve in Fig. 13a, whereas the time warp stretches or shortens the curve along the time axis. We superimpose the volume of the original pairs of data for each subject, as well as the volume evolution computed from the pairwise registration between these pairs of surfaces.

Note that there is no reason that the volume of the registered mean scenario corresponds to the volume evolution computed from the pairwise registration. Indeed, the pairwise registration take into account only a pair of data, whereas the mean scenario integrates the information of the whole database: the mean scenario may contain growth patterns which are not present in a given subject's evolution. Moreover, the pairwise registration aims at minimizing the discrepancy between the two surfaces, whereas the deformed mean scenario is more constrained by the fact that we assume each subject's pair of surfaces to result from a smooth deformation of a mean scenario.

Having said that, we notice that volume evolution of the deformed scenario does not deviate too much from the volume evolution computed from pairwise registration. This means that the morphological deformation accounts well for the different sizes of the structures and that the time warp enables to adjust the slope of the curves to the different growth speed of each subject.

We notice that the curves for one autistic patient are not properly aligned (the patient for which the the volume decreases between the two observations in Fig. 13b). This is the patient detected as an outlier. With the current set of parameters, it was too costly to deform the mean scenario to this subject, which present a unique pattern of size reduction over time (this volume reduction might be real or might be due to a segmentation inaccuracy as well). We run the atlas estimationwith a more important weight for the time warp than for the morphological deformation (σ_ψ/σ_ϕ up to 10 instead of 1) and for larger scales for the time warps ($\lambda_\psi = 1.5$–2.5), which reduces the cost of time warps of large amplitude. We observe that the estimated mean scenario tends to show a volume reduction near age 6 years (after a phase of volume increase from 2 to nearly 6 years) and the outlier is registered to the later part of the mean growth scenario: its time warp shows a strong advance in development of this subject relative to the mean scenario. Nevertheless, this was done at the cost of less accurate registration of all other subjects and the atlas with such parameters was not optimal. Our statistical model prefers to treat this particular subject as an outlier. However, would more subjects be available showing a decreasing volume over time, the atlas would be likely to take this into account by estimating a growth scenario decomposed into a first phase of increasing volume and a second phase of decreasing volume. Then, the subject with decreasing volume would be systematically considered as delayed with respect to the subject with increasing volume at the same age. Such a population, however, would probably violate our main assumption that the growth of the subjects are homologous, in the sense that they derive from a common prototype scenario. It would be better to consider the two sub-groups as two different populations.

4.2 Analysis of the Spatiotemporal Variability

Now, we analyze the spatiotemporal variability of the mean scenario in the population, and not only its effect on the volume distribution. This variability is decomposed into a geometrical part captured by the morphological deformations ϕ^s and a temporal part captured by the time warps ψ^s. Preliminary tests performed on the initial momenta of the 3D diffeomorphisms ϕ^s (see Sect. 3.3) do not show any correlations between the morphological deformations and the class of the subject (autistics, developmental delays and controls). The mean initial momenta of the morphological deformations of each group do not differ significantly from zero. The direction of the first mode of deformation is similar for each group, but the variance is larger for the autistics and developmental delays than for the controls. This mode essentially shows important variations in the elongation of the hippocampus along with an enlargement of the body. By contrast, the depth of the hippocampus almost does not vary.

The time warps are plotted in Fig. 14a for every subject. When the curve is above the $y = x$ axis (resp. below the $y = x$ axis), the evolution of the subject is in advance (resp. is delayed) relatively to the mean scenario. A slope greater than 1 (resp. smaller than 1) denotes an acceleration (resp. a speed reduction) of the evolution of the subject compared to the evolution of the mean scenario. The mean of the curves for each group is plotted in Fig. 14b. Although the mean of

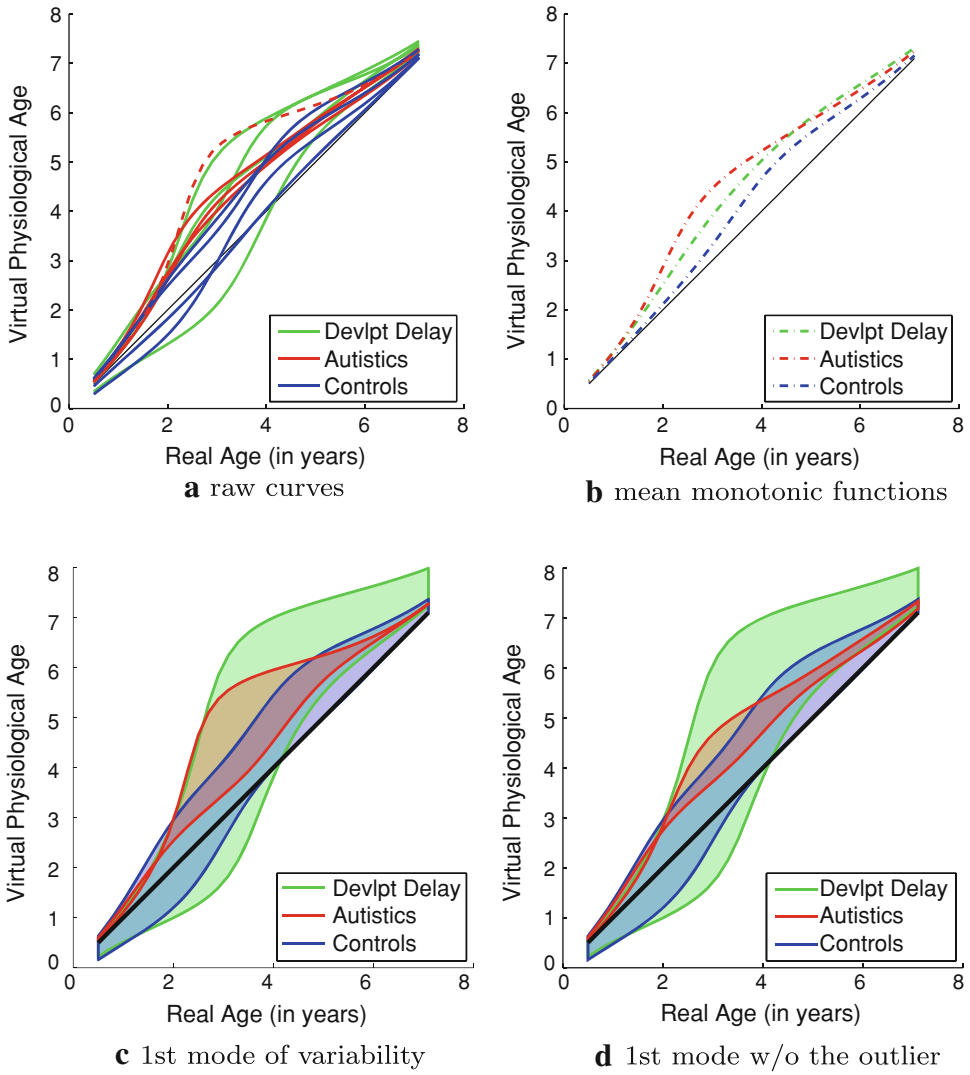


Fig. 14 Estimated time warps from the hippocampus database. (**a**) The *monotonic curves* indicates how the real age of each subject maps to the virtual physiological stage estimated in the mean growth scenario. When *curves* are above the $x = y$ axis, the subject is in advance with respect to the mean scenario. The *dashed red curve* corresponds to the outlier. (**b**) Intrinsic means of each group (also monotonic functions). (**c**) Limits of the first mode of variation at ± 1 standard deviation. (**d**) Same as **c**, but excluding the outlier. It shows that autistics tend to be in advance with respect to the control and that the developmental delays have a much greater variance than the other two groups

all curves seem to be biased (overall over the $x = y$ axis), this bias is not proven to be statistically significant (ratio of mean over standard deviation being equal to 1.21). We compute the intrinsic first mode of variations in the space of smooth monotonic functions (see Sect. 3.3) in Fig. 14c, and in Fig. 14d by excluding the autistic outlier. From these results, we cannot conclude that an autistic or a developmental delayed patient is systematically delayed or in advance compared to controls, even at a given age. However, both the autistics and the controls has a much narrower variability interval than the developmental delays. It seems also that the autistic on average are more advanced than the control at the earlier stages of growth (between 2 and 3 years of age). This suggests that the hippocampus develops faster among autistic children.

To investigate this more in depth, we also compute the spatiotemporal atlas for the amygdala of the right hemisphere (using the the parameters $\lambda_\chi = 15\,\text{mm}$, $\lambda_\phi = 15\,\text{mm}$, $\sigma_\phi = 1$, $\lambda_\psi = 1\,\text{year}$, $\sigma_\psi = 1$, $\lambda_W = 3\,\text{mm}$, $\gamma^\chi = 10^{-3}\,\text{mm}^2$ year, $\gamma^\phi = 10^{-3}\,\text{mm}^2$ and $\gamma^\psi = 10^{-6}\,\text{mm}^4\text{year}^{-2}$). Again, the analysis of the morphological deformations does not highlight informative patterns. By contrast, the analysis of the time warps shown in Fig. 15 reveals that the autistics and the controls share a similar pattern, namely a strong acceleration of the growth with respect to the mean scenario, but at a different age. The acceleration occurs between age 2.5 and 3.5 years for the autistics and between age 4 and age 5 years for the controls. The developmental delays also display a similar pattern, but it occurs at a more variable age. This confirms the hypothesis of an over-growth of the autistics compared to the controls at the earlier stages of development. This also confirms that fact that the developmental delays do not build a very homogeneous group because of much more variable patterns.

These preliminary results on both the geometrical and temporal parts of the variability suggest that the discriminative information between classes might not be inferred from the anatomical variability at a given age, but rather from variations of the growth process. It suggests that autism may more strongly affect the growth speed of the deep brain structures

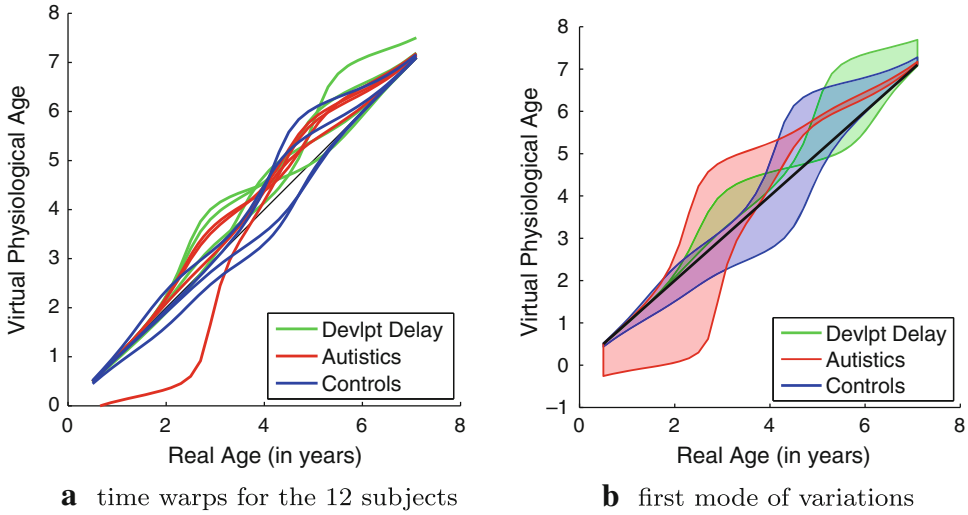


Fig. 15 Estimated time warps from the amygdala database (**a**) time warps for the 12 subjects, (**b**) limits of the first mode of variation at ± 1 standard deviation for each group. Autistics and controls show the same evolution pattern, namely a reduction of speed with respect to the mean scenario (slope smaller than 1) and then a quick acceleration (slope greater than 1). This pattern for the autistics group seems to occur later than for the control group. The developmental delays presents also such pattern but at an arbitrary age. Mean and modes are computed as monotonic functions within the space of 1D diffeomorphisms

rather than its shape, a finding related to brain overgrowth discussed in Hazlett et al. (2011). Note that the hypothesis of an over-growth of the brain of autistic patient has been reported in the literature, for instance in Courchesne et al. (2011). We believe that this new methodology is well adapted to test this hypothesis thanks to the introduction of the time warps, which models explicitly the possible developmental delays between subjects both in shape and in size. Of course, one would need to test it on much larger database: the more time-points per subjects, the more constrained the mean scenario estimation; the more subjects, the more robust the statistics.

5 Comparison of the Endocast Growth Between Chimpanzees and Bonobos

5.1 Framework of the Study

In this section, we aim at using our methodology to characterize the differences of growth patterns between the two closest human relatives: the bonobo (*Pan paniscus*) and the chimpanzee (*Pan troglodytes*). We will also assess the robustness of the method with respect to parameter changes and changes in age labels.

Since bonobos were discovered to science in 1929, the analysis of what distinguishes them from the common chimpanzee has been controversial. After several morphological and behavioral studies (Kuroda 1989; Shea 1989; Kano 1992; de Waal 1995), the hypothesis has emerged that the bonobos may be a "juvenilized" version of the chimpanzee, in the sense that the growth of the bonobos may share common patterns with the one the chimpanzees but with a different tempo. The tools that we have developed, and in particular the introduction of the time warps, seem to be well adapted to test this hypothesis.

For this purpose, we will use one the largest collections of endocasts available for the two species, which comes from the collection of the "Musée de l'Afrique centrale" in Tervuren, Belgium. The endocast is a mold of the endocranium, which provides a replica of the inner surface of the skull and therefore has often played an important role for the analysis of the evolution of the brain in fossil mammals. This data set consists of samples from wild-shot animals: 59 chimpanzees and 60 bonobos, with approximately equal numbers of male and female. They have been scanned with slice thickness between 0.33 and 0.50 mm. The segmentation of the endocasts using `itkSNAP` (Yushkevich et al. 2006) leads to surface meshes. These surfaces have been rigidly co-registered using `gmmreg` (Jian and Vemuri 2005).

It has been observed in Kinzey (1984) that the sequences of teeth emergence in bonobos and chimpanzees are essentially identical. This gives a way to estimate the "dental age" of each skull. We will use this dental age as a common proxy of growth, and not the true age of the specimen, which is not available. As a consequence, each skull has been associated to one the 6 dental ages defined in Shea (1989): infant, child, young juvenile, old juvenile, sub-adult and adult. To refine the classification, we associated some skulls with the intermediate class 'child/young juvenile'. Age distribution is shown in Fig. 1. Without loss of generality, we assume that each dental development stage lasts the same amount of time, namely 1 unit of time. Each unit of time has been discretized with 5 time-points, so that the samples are associated to the time-points $t_i = 5, 10, 15, 20, 25, 30$ according to the dental development of the specimen. The age child/young juvenile has been associated to the time-point 13 (Table 1).

Obviously, this database is cross-sectional by nature. It is unthinkable to have several observations of the same wild

Table 1 Distribution of the dental ages across samples for both species: chimpanzee and bonobo

No. of samples	Infant	Child	Child/young juvenile	Young juvenile	Old juvenile	Sub-adult	Adult	Total
Bonobos	4	8	3	11	7	9	18	60
Chimpanzees	2	6	4	10	13	10	14	59

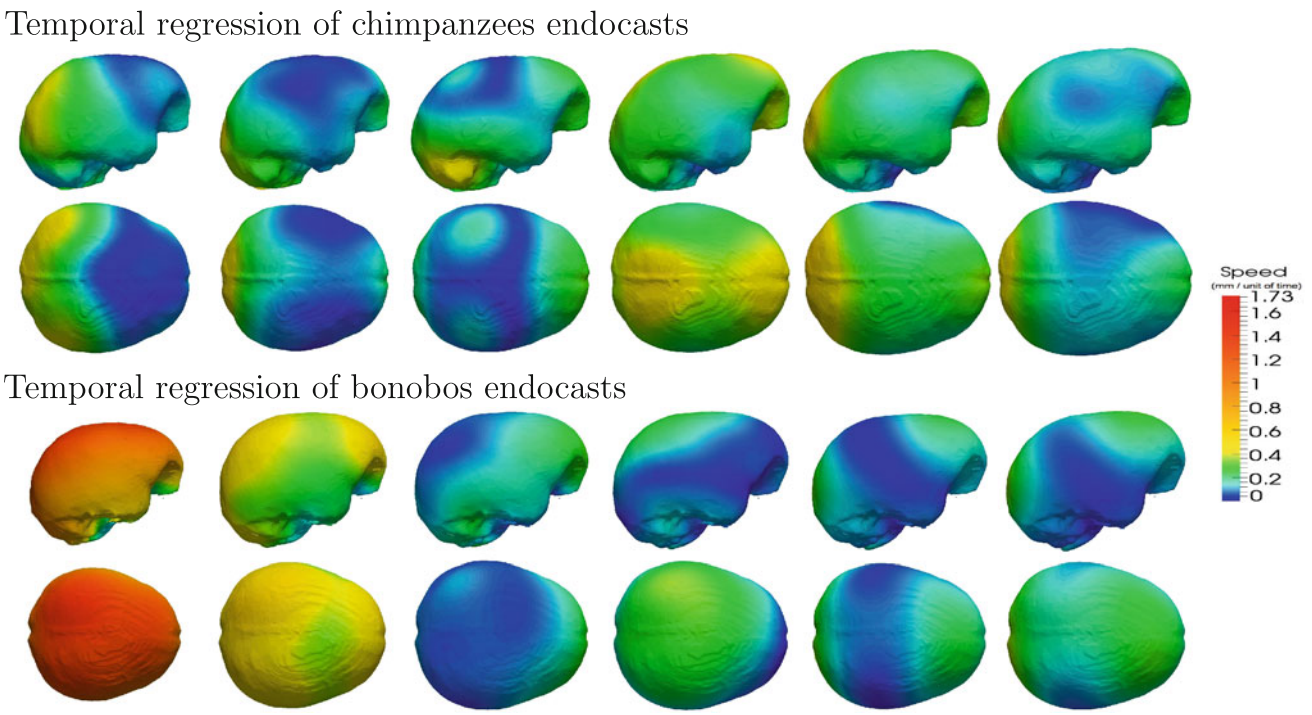


Fig. 16 Temporal shape regression of endocast of the chimpanzees (*top*) and the bonobos (*bottom*) estimated from the original endocasts. In each species, the endocast seems to evolve from a spherical geometry at infancy to an ellipsoidal one at adulthood. However, the dynamics of such changes seem to differ for both species. The quite unrealistic evolution of the chimpanzee endocast at infancy is due to the small amount of data at this age (2). Here, only 6 stages of the growth are shown, although the estimated scenario is *continuous*. Best seen as movies: see Online Resource 2 (chimpanzees) and 3 (bonobos)

animals over time. To make the best of this situation, we choose to estimate first a typical growth scenario for each species independently, applying the regression tool to the cross-sectional data. Second, we analyze the differences between the two growth scenarios using the spatiotemporal registration to measure both morphological differences and possible developmental delays, a key feature we aim at detecting in regards to the bonobos hypothesis.

An alternative approach would consist in using the spatiotemporal atlas construction to estimate an hybrid growth scenario and its deformations to each species, considered as 2 subjects. One the one hand, this would prevent biasing the analysis by choosing a reference species and the inter-species comparison would take into account the fact that the different age groups have a different number of samples. On the other hand, the methodology we choose enables a more direct comparison between species. In particular, the estimation of the species' specific growth scenario, which is done independently for each species, is not constrained by the assumption that the two growth scenarios should derived from the same hybrid scenario.

5.2 Typical Growth Scenario Estimation for Each Species

We choose the smallest endocast within the child class as the baseline S_0 and associate it to the time point $t = 2$. Then, we perform a temporal shape regression of the endocasts (independently for each species), as explained in Sect. 3.1.1 and Algorithm 1. We set the typical spatial interaction between currents $\lambda_W = 10$ mm, the spatial scale of deformation consistency $\lambda^\chi = 20$ mm and the trade-off between fidelity-to-data and regularity $\gamma_\chi = 10^{-3}$ mm^2(unit of time). The diameter of the endocasts are typically between 60 and 70 mm.

Significant samples of the species specific growth scenarios are shown in Fig. 16. The complete scenarios can be seen in the companion movies (see Online Resource 2 (chimpanzees) and 3 (bonobos)). Besides the increase of volume, the most salient effect in both scenarios is an elongation along the posterior/anterior axis and a slight contraction along the superior/inferior axis. As a consequence, the endocast which has an almost spherical geometry at birth has an increasingly ellipsoidal geometry. However, it seems that the chimpanzee endocasts have a stronger anisotropy and that this anisotropy increases faster in time. The subsequent spatiotemporal registration will measure the differences in both scenarios more precisely.

The two growth scenarios differ considerably during infancy and childhood, mainly because of the small amount of data in infancy. The two infant chimpanzees have a larger endocasts compared to both the infant bonobos and the children chimpanzees. To have a more relevant estimation of the growth in infancy, we expect to scan more infant chimpanzees skulls in the future. Note that in the next section we will not take the infancy data into account and will consider the growth scenarios starting at childhood.

We can deduce from the growth scenario an estimation of the evolution of the endocranial volume across ages, as shown in Fig. 17 (Note that we have not performed a regression of the volume but of the shapes instead). Besides the evolution in infancy, one intriguing feature is the apparent decrease in endocranial volume of bonobos at sub-adulthood. This feature is also present in the endocranial volume distribution

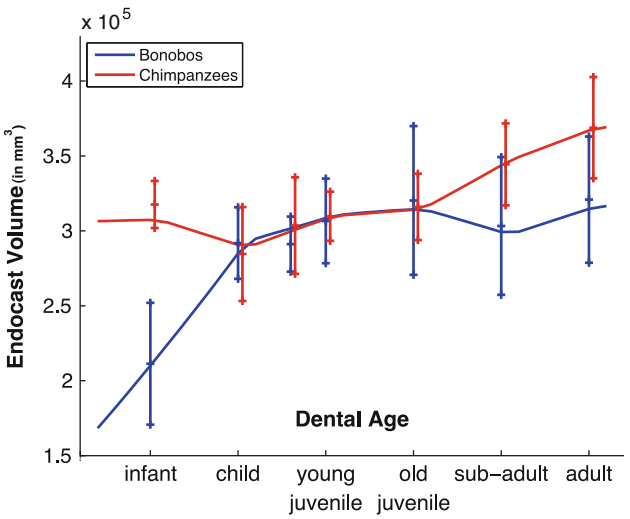


Fig. 17 From the continuous shape regression shown in Fig. 16, we deduce an estimation of the evolution of the endocast volume during growth. Mean and standard deviation of the volume of the original endocasts are superimposed. The intriguing decrease of volume of bonobos at sub-adulthood is not shown to be statistically significant. The unrealistic regression at infancy of chimpanzees is due to the very small number of samples at this age (2)

in the original endocasts (mean and standard deviation are shown in Fig. 17): the mean of the volume at sub-adulthood is smaller than the one of old juveniles. However, the Mann–Whitney U test gives a p value of 0.47 when comparing the volume distribution of old juveniles and sub-adults: the median of the two distributions are not proved to be statistically different. The test run for every pair of consecutive distributions shows a significant increase of volume in only three occasions: (i) between infancy and childhood for the bonobos (p value: 9×10^{-3}); (ii) between childhood and young juvenility for the chimpanzees (p value: 0.07); (iii) and between old-juvenility and sub-adulthood for the chimpanzees (p value: 0.02).

5.3 Spatiotemporal Registration Between the Two Growth Scenarios

We perform a spatiotemporal registration between the two estimated growth scenarios as explained in Sect. 3.1.2 and Algorithm 2. For the reasons explained in the previous section, we consider the part of these scenarios—between childhood and adulthood—discarding the portion between infancy and childhood. We consider the chimpanzee growth scenario as the reference scenario (i.e. the source). The bonobo scenario is sampled every 2 time-steps. These samples play the role of the target shapes. We set the scale of currents to $\lambda_W = 10\,\text{mm}$ as for the regression estimation. We run the registration for different sets of parameters and pick the ones which enable to achieve the smallest discrepancy term in (4). This gives the scale of the morphological deformation: $\lambda^\phi = 10\,\text{mm}$, the scale of the time warp $\lambda^\psi = 1$ unit of time (i.e. duration of one time-point), the spatial power $\sigma_\phi = 40$, the temporal power $\sigma_\psi = 5$, the morphological trade-off $\gamma^\phi = 10^{-5}\,\text{mm}^2$ and the temporal trade-off $\gamma^\psi = 10^{-5}\,\text{mm}^4/(\text{unit of time})$.

The morphological deformation changes the shape of each frame of the chimpanzee growth as shown in Fig. 18 (this is the equivalent figure to the first and second row in Fig. 6, although we plot here an intermediate step of the deformation). It shows that, independently of the age, the bonobos endocasts are rounder than the chimpanzees. The movie of this deformation clearly shows a twist at the anterior and posterior part of the endocast (see companion movie: Online Resource 4).

The graph of the estimated time warp is shown in Fig. 19. This plot shows the correspondence between the ages of the bonobos and the chimpanzees. It mostly shows an important speed reduction of the growth of the bonobos with respect to the chimpanzees between old-juvenility and sub-adulthood. The almost constant slope of the curve during this period of time indicates that the bonobos growth speed is 0.25 times that of the chimpanzees. The graph shows also that the bonobos seem to be slightly in advance with respect to the

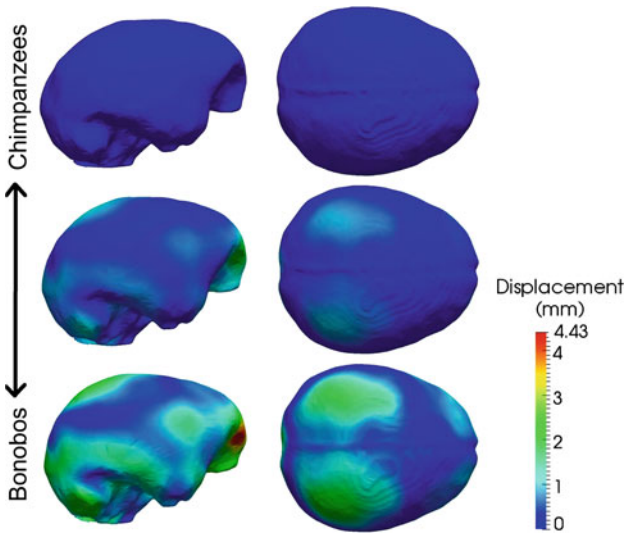


Fig. 18 Morphological part of the spatiotemporal registration between chimpanzees and bonobos growth scenarios. The morphological deformation maps the morphological space of the chimpanzees to that of the bonobos, independently of the age. It is applied here to the chimpanzees endocasts at old juvenility: endocasts from the chimpanzees growth scenario (*top row*), their deformation to the bonobos space (*bottom row*) with an intermediate stage of deformation (*middle row*). This shows that, on average, the endocast of a chimpanzee is more elongated and less round than the one of a bonobo. Note that the deformed endocasts do not match the ones of the bonobo growth at the same age, but at the age given by the correspondence graph shown in Fig. 19. Best seen as a movie: see Online Resource 4

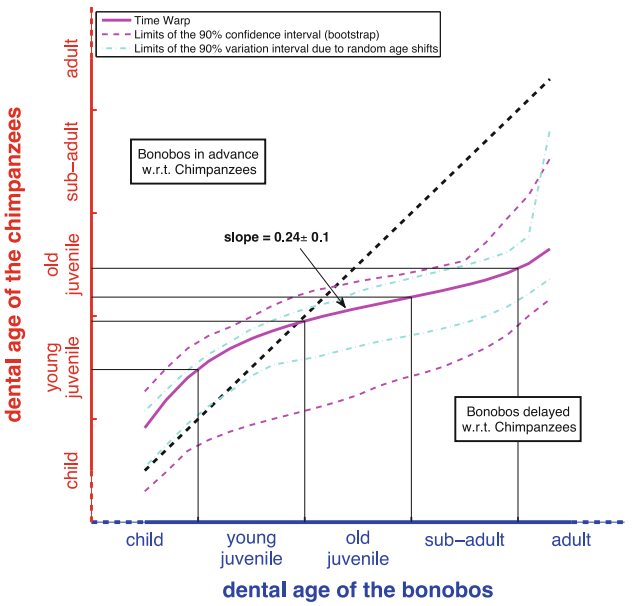


Fig. 19 Time warp between chimpanzee and bonobo growth. It shows that the growth of the bonobos is in advance with respect to the chimpanzees at childhood and then that it drastically slows down during juvenility (almost linearly by a factor 0.24 between old-juvenility and sub-adulthood). This delay seems to decrease at adulthood. *Dashed magenta lines* indicate the limits of the 90 % confidence interval (95 % CI) estimated by bootstrap. *Dashed cyan lines* indicate the limits of the 90 % variation intervals due to random age shifts

chimpanzees at childhood and that the delay of the bonobos growth at sub-adulthood seems to be reduced at adulthood.

We show the effect of this spatiotemporal registration on the evolution of endocast volume of the chimpanzees in Fig. 20 (left). It shows that the spatiotemporal warping enables to match the volume evolution of the chimpanzees endocast closer to that of the bonobos endocast. Besides volume, we also analyze the differences of a measure of the shape, namely the ratio between the height (in the superior-inferior direction) and the width (in the anteroposterior direction) of the endocast. Ratio close to 1 indicates a rounded endocast in the sagittal plane. The evolution of the ratio computed from the species specific scenario is shown in Fig. 20 (right): the decrease of the curves indicates that the endocast become more and more asymmetric (ellipsoidal) during growth. This ratio is always smaller for the chimpanzees than for the bonobos, thus showing a stronger asymmetry for the chimpanzees. As expected, the morphological deformation moves the curves of the chimpanzees closer to the one of the bonobos, except between sub-adulthood and adulthood. Indeed, we have shown in Fig. 18 that the morphological deformation tends give the endocast a more ellipsoidal aspect. The time warp tends to align the slope of the chimpanzee curve to the slope of the bonobos curve. However, we notice that the spatiotemporal warping aligns the evolution of this ratio with less accuracy than that of the volume (see Fig. 20 (left)). It is likely that the registration is primarily driven by the volume information, which may act as a stronger constraint than the shape asymmetries. Multi-scale approaches, with decreasing rigidity scale of the morphological deformation λ_ϕ for instance, should be able to achieve a better matching, which would show an alignment of the measures of shape asymmetries with the same accuracy as the volume.

5.4 Impact of the Temporal Scale

Here, we analyze the variability of the spatiotemporal registration with respect to variations of the temporal scale λ_ψ, while keeping the other parameters fixed. This parameter determines the scale at which the time warp may vary. A large scale means a nearly rigid time warp with very slow variations. A small scale allows quick variations of the time warp during small time intervals.

Figure 21 shows the different values of the data term obtained for different values of the temporal scale λ_ψ. It shows that the optimal value is for $\lambda_\psi = 1$ unit of time, namely the the duration of one age group. This is the value chosen in this study. Figure 22 shows the impact of this temporal scale on the profile of the time warp. The larger the scale, the more rigid the time warp, the less its ability to

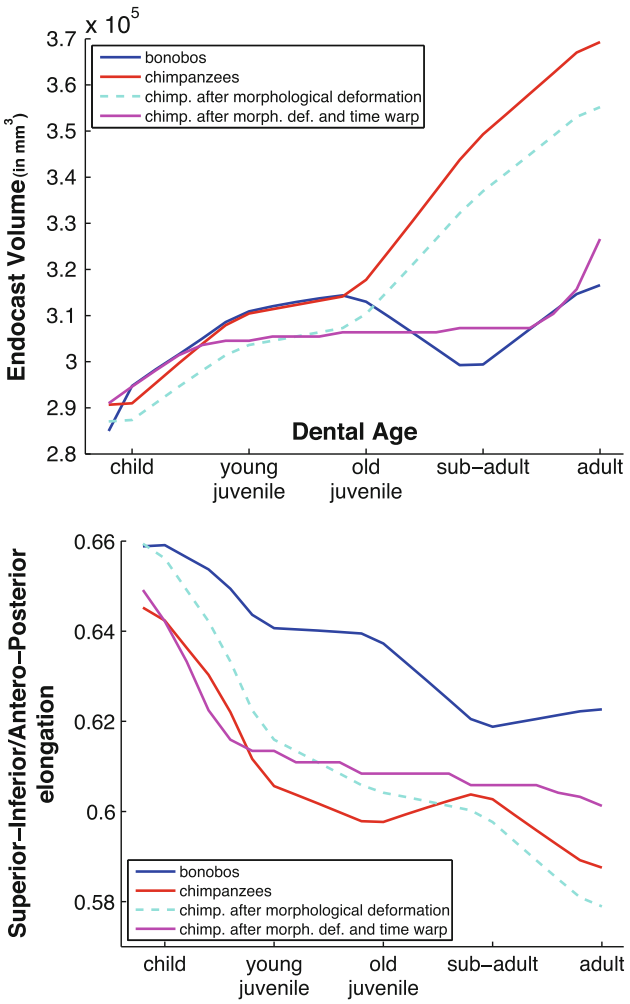


Fig. 20 Effects of the spatiotemporal deformation on the evolution of the volume and the geometry of the endocasts. *Top* Evolution of the endocast volume for the original growth scenarios of both species starting at childhood (in *red* and *blue* as in Fig. 17). *Dashed cyan curve* correspond to the volume of the chimpanzee growth scenario after the morphological deformation. *Magenta curve* is derived from the *cyan curve* by applying the time warp. The combination of the morphological deformation and the time warp approximate the volume evolution of the bonobos. *Bottom* Same experiments but for the evolution of the ratio between the elongation in superior–inferior direction and that in the anteroposterior direction, which gives an indication of how the endocast deviates from a *circular shape* in the sagittal plane. The closer the ratio to 1, the "rounder" the endocast

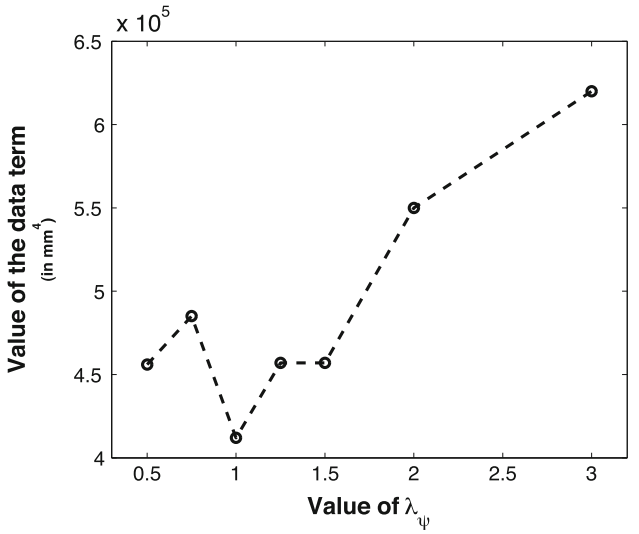


Fig. 21 Effect of the temporal scale λ_ψ on the registration accuracy. Value of the residual data term after registration for different values of the temporal scale λ_ψ (the other parameters being fixed to $\lambda_\phi = 10\,\text{mm}$, $\sigma_\phi = 40$, $\sigma_\psi = 5$, $\gamma_\phi = 10^{-5}$ and $\gamma_\psi = 10^{-5}$). It indicates the optimal value of $\lambda_\psi = 1$ unit of time

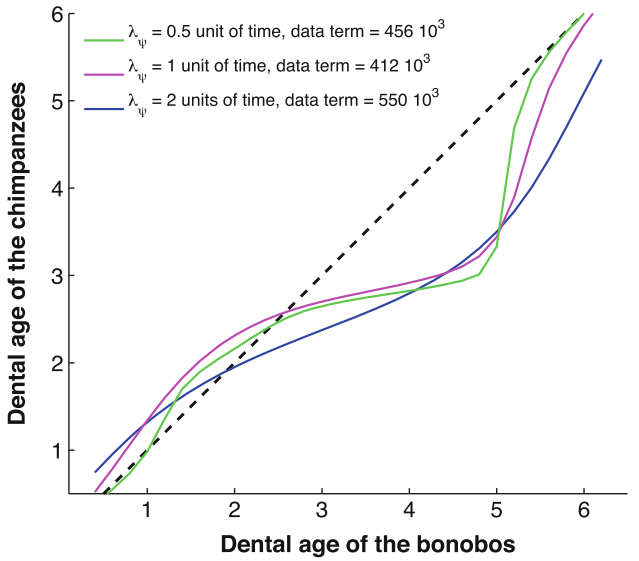


Fig. 22 Effect of the temporal scale λ_ψ on the time warp. If λ_ψ is too large, it cannot capture fast variations in the dynamics of growth of both species. If λ_ψ is too small, it costs more to capture the large-scaled variations. Optimal solution is for $\lambda_\psi = 1$ unit of time

capture highly non-linear variations. The smaller the scale, the more expensive the cost of a regular deformation.

5.5 Estimation of Confidence Intervals via Bootstrap

In this section, we aim at studying the robustness of the estimation of the growth scenario with respect to the samples we have. We use here a bootstrap procedure: a resampling with replacement procedure is applied within each age group of each species, yielding a new data set of the same size as the original. Then, we estimate the two growth scenarios using this new set of data. We repeat the procedure 100 times, so that we end up with 100 growth scenarios for each species, which simulates the variability of these scenarios with respect to the choice of the samples.

For each simulation, we compute the evolution of the endocranial volume given by the estimated growth scenario. Fig 23 (top, left) shows the 95 % CI of the volume evolutions. It shows that the variability of the volume estimation

is relatively small, in particular with respect to the difference in volume between the two species at adulthood.

To gain more insight into the variability of the geometry of the growth scenario, we compute the distance between the endocast of a simulated growth scenario and the reference one at each time-step, using the norm on currents. In Fig. 23 (top,right), we show the 95 % CI of this difference expressed in terms of percentage of the norm of the endocast of the reference scenario. It shows a large deviation in bonobos infancy and chimpanzees childhood. This large variability is expected due to the small number and the large variability of data at those ages. This confirms that more data are needed for a more robust estimation of the endocranial growth at these ages. The variance of the endocranial geometry during old-juvenility and sub-adulthood is much larger for the bonobos than for the chimpanzees. This may explain the decrease of volume at bonobos sub-adulthood, as many more samples would be needed to converge to the "true" mean. This might also indicate a bi-modal distribution for male and female.

Eventually, we notice that the estimation of the volume seems much less variable than the estimation of the geometry. This is not surprising since one needs much more data to robustly estimate the whole geometry (which has potentially an infinite number of degrees of freedom) than the scalar measure of volume.

Then, we compute the spatiotemporal registration between every pair of growth scenarios. The estimation of the 95 % CI of the time warp is shown in Fig. 19. These experiments allow us to give also a confidence interval of the developmental delay between the bonobos and the chimpanzees: the slope of the time warp at old juvenility falls into the interval [0.14, 0.34] in 90 % of the cases.

5.6 Estimation of Variability Intervals via Random Age Shifts

Here, we study the robustness of our estimations to perturbations of the age estimates. For this purpose, we simulate age perturbations by adding a zero-mean Gaussian variable with standard deviation 1 time-point to the dental age of each sample. This means that in 50 % of the cases the dental ages have been shifted by $+1$ or -1 time-point, in 10 % of the cases they have been shifted by more than one time-point, and in 40 % of the cases they have not moved. We recall that the duration of every dental age group was of five time-points (i.e. 1 unit of time) in the original experiments, which means that in 10 % of the cases, the age estimate was shifted at or beyond the boundaries its group. Given these new age estimates, we compute two growth scenarios and then the spatiotemporal deformations between the portion of the scenarios between childhood and adulthood. We repeat this procedure 100 times. We define a 90 % variability interval by discarding the five largest and five smallest values of any scalar measurements taken out of these simulations.

In Fig. 23 (bottom,left), we show the limits of the 90 % variability interval of the volume evolution measured from the growth scenario. In Fig. 23 (bottom,right), we show the 90 % variability interval of the distance between the perturbed growth scenario and the reference one.

We compute the spatiotemporal registration between every pair of perturbed growth scenarios. The 90 % variability interval of the estimated time warp is shown in Fig. 19.

All these experiments show that the perturbations induced by randomly shifting the age estimates are always smaller than the perturbations due to the bootstrap resampling. Therefore, these perturbations are not statistically significant, thus illustrating the robustness of our method to the uncertainty of the age estimates.

6 Conclusion

In this paper, we introduced a new growth model for shape evolution, as a continuous diffeomorphic deformation of a baseline shape over time. Its estimation allows us to infer a continuous shape trajectory from few observations sparsely distributed in time.

We introduced a new registration scheme for the comparison of individual growth trajectories. We proposed to model the differences between individual trajectories as a combination of time-independent shape differences and a difference in the pace of shape changes over time. For this purpose, we introduced the generic construction of 1D diffeomorphisms: smooth and monotonic functions called time warps. This model supposes that the individual trajectories are comparable, in the sense that they contain the same growth patterns, but with a different appearance and a different timing. The presence of new growth patterns in one of the trajectories would not be considered by the registration and would lead to a large residual misalignment after the registration.

Eventually, we introduced an original statistical model for the study of longitudinal databases. This model assumes that the consecutive observations of different subjects derive from a spatiotemporal deformation of an unknown prototype scenario of evolution. The estimation of this models leads to an atlas, which consists of (i) a prototype shape called the template, (ii) the continuous diffeomorphic deformation of this prototype shape over time, called the mean scenario of evolution and (iii) the spatiotemporal deformations of this mean scenario, which map the set of time-dependent observations of each subject. The template shape and the mean scenario of evolution captures the invariants in the population: common shape features and common growth patterns. The spatiotemporal deformations summarize the variations

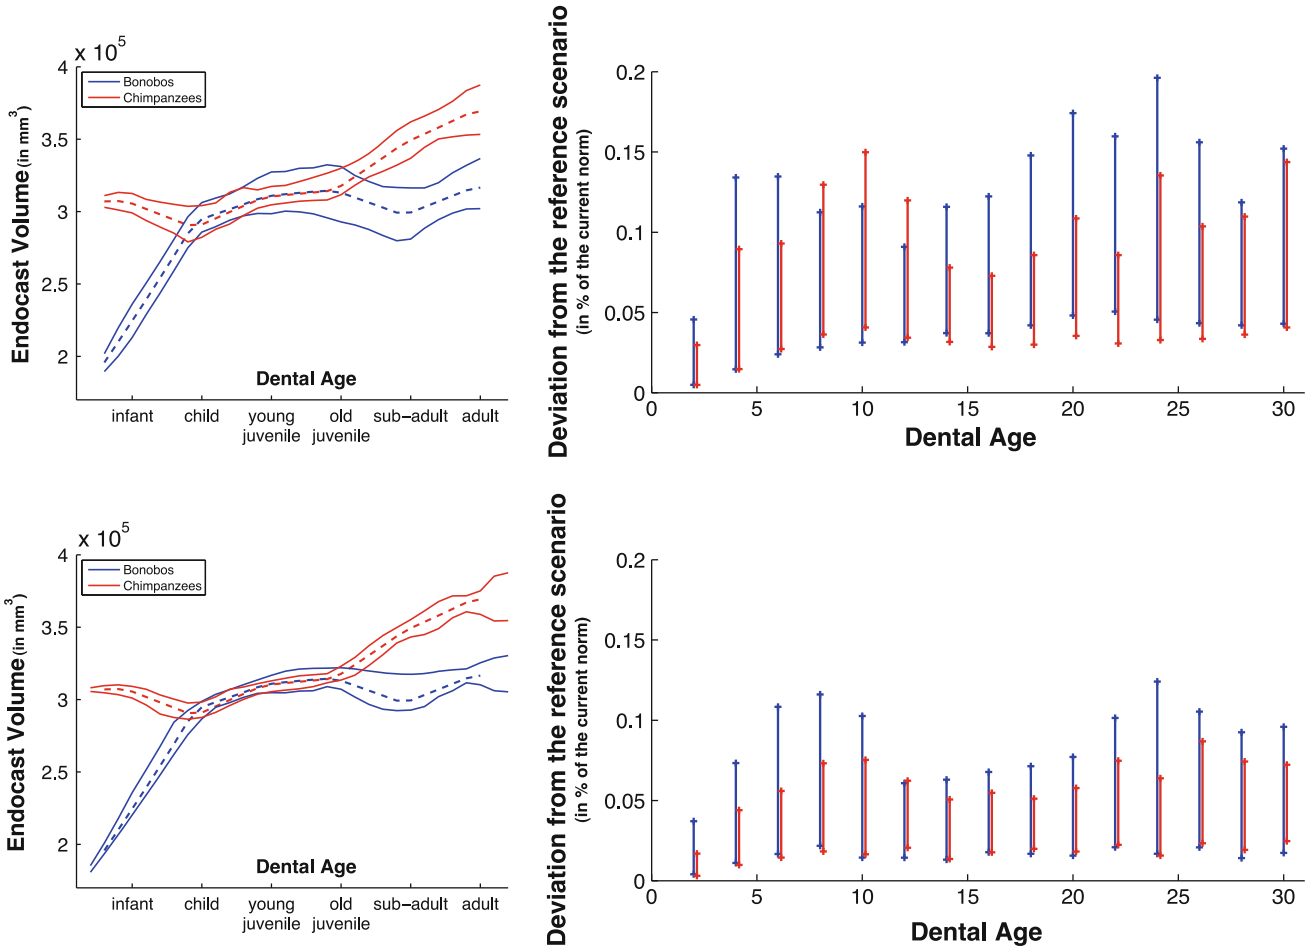


Fig. 23 Bootstrap confidence intervals due to resampling (*top*) and due to random age shifts (*bottom*). *Left* Evolution of the endocranial volume given by the reference growth scenario (*bold line*) and its 95 % CI estimated via a bootstrap procedure. *Right* Discrepancy between the reference growth scenario and the ones estimated by bootstrap, measured as the current norm between the frames. 90 % confidence interval is shown at every time-step. On average, the bootstrap makes the frames to vary in the space of currents within a neighborhood of radius 10 % the norm of the reference frames

of these invariants in terms of various appearances and various paces of shape change. The characteristics, which are specific to a given individual, are not taken into account and are considered as outliers since it does not comply with the hypotheses of the model.

In addition, we provided a statistical framework to compute intrinsic statistics in the space of spatiotemporal deformations (space of 3D diffeomorphisms for the morphological deformation and space of smooth monotonic functions for the time warp). Mean and modes of deformations are computed as deformations themselves. This is not only useful to perform statistical processing like hypothesis testing or classification for instance, but also gives a way to visualize and interpret the geometrical features captured by the model. Our method could be seen as a longitudinal extension of the concept of Karcher mean, which estimates at the same time the mean (first moment) and the covariance (second moment) from data lying on a manifold. Extensions of other manifold-related methods for shape averaging could be also investigated (Gerber et al. 2010; Xie et al. 2010)

We illustrated our approach with two biological examples. The study of a longitudinal database of deep brain structures from autistic, developmental delays, and controls, shows that a given pathology might be characterized more by the differences in the pace of maturation of a structure, rather than the differences in shapes at a given age. The study of a time-series data set of endocast of bonobos and chimpanzees allows us to give new insights into the differences in terms of development of the endocranium between these genetically very close species. In particular, we were able to give an estimate of the expected developmental delay of bonobos relative to chimpanzees at juvenility. Moreover, we show the robustness of this estimation with respect to variations in the age estimates.

The principle approach that we proposed leads to the definition of new objective functions, for which we proposed efficient algorithms. The shape regression and the spatiotemporal registration are solved using a single gradient descent scheme. The estimation of the spatiotemporal atlas relies on an alternated minimization procedure. Future work will investigate the possibility to derive a single gradient scheme for the estimation of the whole atlas. This should lead to a faster and more controllable procedure. Another important direction of research is to include the residuals into the statistical analysis in the spirit of Durrleman et al. (2009a, 2011), in order to better characterize possible outliers, or to detect consistent subgroups in the population.

We emphasized that our approach relies on a set of explicit hypotheses and that alternatives models are also worth investigating. We hope that this approach, which has been driven more by methodological considerations than by the applications, will help to better appreciate the challenges, which are inherent to the joint modeling of time and shape variations. This might contribute to the emergence of an axiomatic approach for the statistical analysis of longitudinal shape data.

Acknowledgments We would like to thank B. Combès (IRISA, France) for preprocessing the endocast data, J. Piven, director of Carolina Institute for Developmental Disabilities at UNC Chapel Hill, for providing imaging data related to autism research, and M. Styner (Psychiatry UNC Chapel Hill) for processing the subcortical structures. We thank W. Van Neer and E. Gilissen the previous and current curator of the "Musée de l'Afrique Centrale" at Tervuren (Belgium). We are indebted to Chems Touati for his help for creating figures and movies and James Fishbaugh for his kind proofreading of the manuscript, both at the Scientific Computing and Imaging Institute, University of Utah. This work has been funded in part by the INRIA ARC `3D-Morphine` (PI: Sylvain Prima), the European IP Project `Health-e-child` (IST-2004-027749) and Microsoft Research.

A Differentiation of the Temporal Shape Regression Criterion

A.1 Matrix Notations

For the sake of simplicity, we introduce matrix notations: $\mathbf{x}_0 = \{x_p\}_{p=1,\dots,N}$ denotes the $3N$ vector which is the concatenation of the coordinates of N vertices of the baseline shape M_0. Using the notations of Sect. 3.2.2, we denote the moving points $\mathbf{x}(t)$ (resp. the parameterizing vectors $\boldsymbol{\alpha}(t)$) the $3N$ vector: $(x_p(t))_{p=1,\dots,N}$ (resp. $(\alpha_p(t))_{p=1,\dots,N}$). We denote also $\mathbf{K}^\chi(\mathbf{x}(t),\mathbf{x}(t))$ the $3N$-by-$3N$ block matrix whose block p,q is given by the 3-by-3 matrix $(K^\chi(x_p(t),x_q(t)))$. This matrix is symmetric, positive definite by definition of the kernel K^χ.

Thanks to these notations, the norm of the speed vector v_t is written: $\|v_t\|^2_{V^\chi} = \boldsymbol{\alpha}(t)^t\mathbf{K}^\chi(\mathbf{x}(t),\mathbf{x}(t))\boldsymbol{\alpha}(t)$. For A, a function from $\mathbb{R}^3$ to $\mathbb{R}$, we denote by d_xA its Jacobian matrix at point x, so that for any vector V: $(d_xA)V = (\nabla_xA)^tV$. By extension, $\nabla_{\mathbf{x}}A$ denotes the $3N$ vector $(\nabla_{x_1}A,\dots,\nabla_{x_N}A)$.

With these notations, the regression criterion (17) becomes:

$$E\left((\boldsymbol{\alpha}(t))_{t\in[0,T]}\right) = \sum_{t_i} A_i(\mathbf{x}_{t_i}) + \int_0^T L^\chi(\mathbf{x}(t),\boldsymbol{\alpha}(t))dt \quad (39)$$

subject that:

$$\frac{d\mathbf{x}(t)}{dt} = f(\mathbf{x}(t),\boldsymbol{\alpha}(t)) \quad \text{with} \quad \mathbf{x}(0) = \mathbf{x}_0 \quad (40)$$

where we denote:

$$\begin{aligned} f(\mathbf{x}(t),\boldsymbol{\alpha}(t)) &= \mathbf{K}^\chi(\mathbf{x}(t),\mathbf{x}(t))\boldsymbol{\alpha}(t) \\ L^\chi(\mathbf{x}(t),\boldsymbol{\alpha}(t)) &= \gamma_\chi\boldsymbol{\alpha}(t)^t f(\mathbf{x}(t),\boldsymbol{\alpha}(t)) \end{aligned} \quad (41)$$

For the sake of simplicity, we will denote in the sequel, $f(t)$ and $L^\chi(t)$ instead of $f(\mathbf{x}(t),\boldsymbol{\alpha}(t))$ and $L^\chi(\mathbf{x}(t),\boldsymbol{\alpha}(t))$.

A.2 Gradient in a matrix form

Let δE be a variation of the criterion E with respect to a variation $\delta\boldsymbol{\alpha}(t)$ of the momenta $\boldsymbol{\alpha}(t)$:

$$\begin{aligned} \delta E = \sum_{t_i}\left(d_{\mathbf{x}(t_i)}A_i\right)\delta\mathbf{x}(t_i) + \int_0^T (\partial_{\mathbf{x}}L^\chi(t))\delta\mathbf{x}(t) \\ +(\partial_{\boldsymbol{\alpha}}L^\chi(t))\delta\boldsymbol{\alpha}(t)dt \end{aligned} \quad (42)$$

where $\delta\mathbf{x}(t)$ denotes the variations of the positions $\mathbf{x}(t)$ with respect to the variations of the momenta $\boldsymbol{\alpha}(t)$. The differentiation of the flow Eq. (40) shows that these variations $\delta\mathbf{x}(t)$ satisfy a linear ODE with source term:

$$\begin{aligned} \frac{d}{dt}\delta\mathbf{x}(t) = (\partial_{\mathbf{x}}f(t))\delta\mathbf{x}(t) \\ +(\partial_{\boldsymbol{\alpha}}f(t))\delta\boldsymbol{\alpha}(t) \quad \text{with} \quad \delta\mathbf{x}(0) = 0 \end{aligned} \quad (43)$$

We introduce the flow R_{ut} for $u,t\in[0,T]$ which is solution of the homogeneous equation:

$$\frac{dR_{ut}}{dt} = R_{ut}(\partial_{\mathbf{x}}f(t)) \quad \text{with } R_{tt} = \mathrm{Id} \quad (44)$$

The method of the variations of the parameters leads to the following solution of the ODE:

$$\delta\mathbf{x}(t) = \int_0^t R_{ut}\partial_{\boldsymbol{\alpha}}f(u)\delta\boldsymbol{\alpha}(u)du \quad (45)$$

In particular, this shows that we can write the variations $\delta\mathbf{x}(t_i)$ as:

$$\delta\mathbf{x}(t_i) = \int_0^T R_{tt_i}\partial_{\boldsymbol{\alpha}}f(t)\delta\boldsymbol{\alpha}(t)\mathbf{1}_{\{t\le t_i\}}dt \quad (46)$$

where $\mathbf{1}_{\{t\le t_i\}} = 1$ if $t\le t_i$ and 0 otherwise.

Now, we can plug these last two equations into (42). Using Fubini's theorem, which implies that $\int_0^T \int_0^t F(u,t)dudt = \int_0^T \int_u^T F(u,t)dtdu = \int_0^T \int_t^T F(t,u)dudt$ for any L^2 function $F(u,t)$, this leads to:

$$\delta E = \int_0^T \Big(\partial_{\boldsymbol{\alpha}} L^\chi(t) + \underbrace{\Big(\sum_i d_{\mathbf{x}(t_i)} A_i R_{tt_i} \mathbf{1}_{\{t \le t_i\}} + \int_t^T \partial_{\mathbf{x}} L^\chi(u) R_{tu} du \Big)}_{\eta(t)^t} \partial_{\boldsymbol{\alpha}} f(t) \Big) \delta\boldsymbol{\alpha}(t) dt \tag{47}$$

This gives the gradient of E with respect to the L^2 metric as:

$$\nabla_{\boldsymbol{\alpha}} E(t) = \partial_{\boldsymbol{\alpha}} L^\chi(t)^t + \partial_{\boldsymbol{\alpha}} f(t)^t \boldsymbol{\eta}(t) \tag{48}$$

where we denote the auxiliary variable $\boldsymbol{\eta}(t)$:

$$\boldsymbol{\eta}(t) = \sum_i (R_{tt_i})^t \nabla_{\mathbf{x}(t_i)} A_i \mathbf{1}_{\{t \le t_i\}} + \int_t^T (R_{tu})^t \partial_{\mathbf{x}} L^\chi(u)^t du \tag{49}$$

The auxiliary variable $\boldsymbol{\eta}(u)$ depends on the flows R_{ut} and therefore satisfies an ODE. To make this ODE explicit, we write the inverse flow R_{ut} in integral form. Noticing that $R_{tu} R_{ut} = \mathrm{Id}$, we have $\frac{dR_{ut}}{du} = -\partial_{\mathbf{x}} f(u) R_{ut}$, which gives in integral form (noticing that R_{ut} and f commute):

$$R_{ut} = \mathrm{Id} + \int_u^t R_{st} \partial_{\mathbf{x}} f(s) ds. \tag{50}$$

Now, we can plug this equation into the definition of $\boldsymbol{\eta}(t)$ in (49). Writing $R_{tt_i} = \mathrm{Id} + \int_t^T R_{ut_i} \partial_{\mathbf{x}} f(u) \mathbf{1}_{\{u \le t_i\}} du$ and noticing that for any L^2 function $F(u,s)$, the Fubini's theorem implies that $\int_t^T \int_t^u F(u,s)dsdu = \int_t^T \int_u^T F(s,u)dsdu$, this leads to:

$$\boldsymbol{\eta}(t) = \sum_i \nabla_{\mathbf{x}(t_i)} A_i \mathbf{1}_{\{t \le t_i\}} + \int_t^T \partial_{\mathbf{x}} L^\chi(u)^t + \underbrace{\partial_{\mathbf{x}} f(u)^t \left(\sum_i (R_{ut_i})^t \nabla_{\mathbf{x}(t_i)} A_i \mathbf{1}_{\{u \le t_i\}} \mathbf{1}_{\{t \le t_i\}} + \int_u^T (R_{us})^t \partial_{\mathbf{x}} L^\chi(s)^t ds \right)}_{(\star)} du \tag{51}$$

Now, we notice that $t \le u$ within the integral, which implies that $\mathbf{1}_{\{t \le t_i\}} \mathbf{1}_{\{u \le t_i\}} = \mathbf{1}_{\{u \le t_i\}}$. Hence, $(\star)$ is exactly equal to $\boldsymbol{\eta}_u$. Therefore, $\boldsymbol{\eta}_t$ is the solution of the integral equation (integrated upstream in time):

$$\boldsymbol{\eta}(t) = \sum_i \nabla_{\mathbf{x}_{t_i}} A_i \mathbf{1}_{\{t \le t_i\}} + \int_t^T \partial_{\mathbf{x}} L^\chi(u)^t + \partial_{\mathbf{x}} f(u)^t \boldsymbol{\eta}(u) du \tag{52}$$

A.3 Gradient in Coordinates

Due to the definition of the functions f and L^χ in (41), we have:

$$\begin{aligned} \partial_{\mathbf{x}} f &= (\partial_1 + \partial_2)(\mathbf{K}^\chi(\mathbf{x},\mathbf{x})\boldsymbol{\alpha}) & \partial_{\boldsymbol{\alpha}} f &= \mathbf{K}^\chi(\mathbf{x},\mathbf{x}) \\ \partial_{\mathbf{x}} L^\chi &= \gamma_\chi \boldsymbol{\alpha}^t \left((\partial_1 + \partial_2) \mathbf{K}^\chi(\mathbf{x},\mathbf{x})\boldsymbol{\alpha} \right) & \partial_{\boldsymbol{\alpha}} L^\chi &= 2\gamma_\chi \boldsymbol{\alpha}^t \mathbf{K}^\chi(\mathbf{x},\mathbf{x}) \end{aligned} \tag{53}$$

Therefore, the gradient of the regression criterion with respect to the L^2 metric given in (48) is now equal to:

$$\nabla_{\boldsymbol{\alpha}} E(t) = \mathbf{K}^\chi(\mathbf{x}(t),\mathbf{x}(t)) \left(2\gamma_\chi \boldsymbol{\alpha}(t) + \boldsymbol{\eta}(t) \right).$$

The matrix $\mathbf{K}^\chi(\mathbf{x}(t),\mathbf{x}(t))$ is precisely the Sobolev metric induced by the kernel on the set of L^2 functions (see Sect. 3.2.2), so that the gradient with respect to this metric is given by:

$$\nabla_{\alpha_p} E(t) = 2\gamma_\chi \alpha_p(t) + \eta_p(t) \tag{54}$$

The auxiliary variable $\boldsymbol{\eta}(t)$ satisfies the ODE (52), now written as:

$$\boldsymbol{\eta}(t) = \sum_i \nabla_{\mathbf{x}_{t_i}} A_i \mathbf{1}_{\{t \le t_i\}} + \int_t^T \Big((\partial_1 + \partial_2) \mathbf{K}^\chi(\mathbf{x}(u), \mathbf{x}(u))\boldsymbol{\alpha}(u) \Big)^t \left(\gamma_\chi \boldsymbol{\alpha}(u) + \boldsymbol{\eta}(u) \right) du \tag{55}$$

The $3N$ vector $\nabla_{\mathbf{x}_{t_i}} A_i$ is equal to $(\nabla_{x_1(t_i)} A_i, \ldots, \nabla_{x_N(t_i)} A_i)$. For generic $3N$ vectors $\mathbf{x}$, $\mathbf{y}$ and $\boldsymbol{\alpha}$, the kth coordinate of the $3N$-vector $\mathbf{K}^\chi(\mathbf{x},\mathbf{y})\boldsymbol{\alpha}$ is given as: $(\mathbf{K}^\chi(\mathbf{x},\mathbf{y})\boldsymbol{\alpha})_k = \sum_{p=1}^N K^\chi(x_k, y_p)\alpha_p$. The kernel K^χ is scalar, namely of the form $K^\chi(x,y) = k^\chi(x,y)\mathrm{Id}$ for a scalar function k^χ. We have therefore for every $i, j = 1, \ldots, N$:

$$\begin{aligned} \partial_{x_i} \left(\mathbf{K}^\chi(\mathbf{x},\mathbf{y})\boldsymbol{\alpha} \right)_j &= \sum_{p=1}^N \alpha_p \left(\nabla_1 k^\chi(x_i, y_p) \right)^t \delta(i-j) \\ \partial_{y_i} \left(\mathbf{K}^\chi(\mathbf{x},\mathbf{y})\boldsymbol{\alpha} \right)_j &= \alpha_i \left(\nabla_2 k^\chi(x_j, y_i) \right)^t \end{aligned} \tag{56}$$

Therefore, for a generic $3N$-vector $\boldsymbol{\beta}$, we have:

$$\Big((\partial_1 + \partial_2) \left(\mathbf{K}^\chi(\mathbf{x},\mathbf{y})\boldsymbol{\alpha} \right)^t \boldsymbol{\beta} \Big)_k = \sum_{p=1}^N \alpha_p^t \beta_k \nabla_1 k^\chi(x_k, y_p) + \alpha_k^t \beta_p \nabla_2 k^\chi(x_p, y_k) \tag{57}$$

Now, we can apply this equation with $\mathbf{y} = \mathbf{x}$ and $\boldsymbol{\beta} = \gamma_\chi \boldsymbol{\alpha} + \boldsymbol{\eta}$ and combine it with (55). Noticing that for a symmetric kernel, we have $\nabla_1 k(x,y) = \nabla_2 k(y,x)$, we get eventually

the set of ODEs satisfied by the functions $\eta_p(t)$ as given in (19).

B Differentiation of the Spatiotemporal Matching Criterion

B.1 Matrix Notations

Let $\mathbf{t}_0 = \{t_j\}_{j=1,\dots,N_{\text{target}}}$ be the vector of time-points associated to the target shapes. The 1D diffeomorphism ψ_u changes $\mathbf{t}_0$ into $\mathbf{t}(u) = \{t_j(u)\}_{j=1,\dots,N_{\text{target}}}$ for $u \in [0,1]$. This vector satisfies the ODE: $\frac{d\mathbf{t}}{du}(u) = \mathbf{K}^{\psi}(\boldsymbol{\beta}(u), \boldsymbol{\beta}(u))\mathbf{t}(u)$ with $\mathbf{t}(0) = \mathbf{t}_0$, where $\boldsymbol{\beta}(u)$ is the concatenation of the vectors $\beta_j(u)$ defined in (23), $\mathbf{K}^{\psi}(\boldsymbol{\beta}(u), \boldsymbol{\beta}(u))$ is the block matrix whose block (i, j) is given by: $K^{\psi}(\beta_i(u), \beta_j(u))$.

Similarly, we denote $\mathbf{x}_0(t) = \{x_p(t)\}_{p=1,\dots,N}$ be the concatenation of the positions of all the points of the source evolution $S(t)$ for any time-point t and $\mathbf{x}_0(\mathbf{t}(1))$ the concatenation of the $\mathbf{x}_0(t_j(1))$ for $j = 1, \dots, N_{\text{target}}$. The diffeomorphism ϕ_u maps this vector to $\mathbf{x}(u)$, which satisfies the ODE: $\frac{d\mathbf{x}}{du} = \mathbf{K}^{\phi}(\mathbf{x}(u), \mathbf{x}(u))\boldsymbol{\alpha}(u)$ with initial condition: $\mathbf{x}(0) = \mathbf{x}_0(\mathbf{t}(1))$ (which depends on the final time-points $\mathbf{t}(1)$), where $\boldsymbol{\alpha}(u)$ is the concatenation of the vectors $\alpha_{p,j}(u)$ defined in (22) for $p = 1, \dots, N$ and $j = 1, \dots, N_{\text{target}}$.

Therefore, we can write the matching criterion (4) as:

$$E(\boldsymbol{\alpha}(u), \boldsymbol{\beta}(u)) = A(\mathbf{x}(1)) + \int_0^1 L^{\phi}(\mathbf{x}(u), \boldsymbol{\alpha}(u))\, du + \int_0^1 L^{\psi}(\mathbf{t}(u), \boldsymbol{\beta}(u))\, du \tag{58}$$

subject to:

$$\begin{cases} \dfrac{d\mathbf{x}(u)}{du} = f(\mathbf{x}(u), \boldsymbol{\alpha}(u)) \quad \text{with} \quad \mathbf{x}(0) = \mathbf{x}_0(\mathbf{t}(1)) \\ \dfrac{d\mathbf{t}(u)}{du} = g(\mathbf{t}(u), \boldsymbol{\beta}(u)) \quad \text{with} \quad \mathbf{t}(0) = \mathbf{t}_0 \end{cases} \tag{59}$$

where we denote:

$$\begin{aligned} f(\mathbf{x}(u), \boldsymbol{\alpha}(u)) &= \mathbf{K}^{\phi}(\mathbf{x}(u), \mathbf{x}(u))\boldsymbol{\alpha}(u) \\ g(\mathbf{t}(u), \boldsymbol{\beta}(u)) &= \mathbf{K}^{\psi}(\mathbf{t}(u), \mathbf{t}(u))\boldsymbol{\beta}(u) \\ L^{\phi}(\mathbf{x}(u), \boldsymbol{\alpha}(u)) &= \gamma_{\phi}\boldsymbol{\alpha}(u)^t f(\mathbf{x}(u), \boldsymbol{\alpha}(u)) \\ L^{\psi}(\mathbf{t}(u), \boldsymbol{\beta}(u)) &= \gamma_{\psi}\boldsymbol{\beta}(u)^t g(\mathbf{t}(u), \boldsymbol{\beta}(u)) \end{aligned} \tag{60}$$

For the sake of simplicity, we will write in the sequel $f(u)$, $g(u)$, $L^{\phi}(u)$ and $L^{\psi}(u)$ instead of $f(\mathbf{x}(u), \boldsymbol{\alpha}(u))$, $g(\mathbf{t}(u), \boldsymbol{\beta}(u))$, $L^{\phi}(\mathbf{x}(u), \boldsymbol{\alpha}(u))$ and $L^{\psi}(\mathbf{t}(u), \boldsymbol{\beta}(u))$ respectively.

B.2 Gradient in a Matrix Form

Now, let δE be a variation of the criterion E induced by a variation of the momenta $\delta\boldsymbol{\alpha}(u)$ and $\delta\boldsymbol{\beta}(u)$:

$$\begin{aligned} \delta E = {} & (d_{\mathbf{x}(1)}A)\delta\mathbf{x}(1) + \int_0^1 (\partial_{\mathbf{x}}L^{\phi}(u))\delta\mathbf{x}(u) \\ & + (\partial_{\boldsymbol{\alpha}}L^{\phi}(u))\delta\boldsymbol{\alpha}(u) + (\partial_{\mathbf{t}}L^{\psi}(u))\delta\mathbf{t}(u) \\ & + (\partial_{\boldsymbol{\beta}}L^{\psi}(u))\delta\boldsymbol{\beta}(u)du \end{aligned} \tag{61}$$

where we denote $\delta\mathbf{x}(u)$ and $\delta\mathbf{t}(u)$ the variations of the path $\mathbf{x}(u)$ and $\mathbf{t}(u)$ induced by the variations of the momenta $\delta\boldsymbol{\alpha}(u)$ and $\delta\boldsymbol{\beta}(u)$. These vectors satisfy the linear ODEs with source term derived from (59):

$$\begin{aligned} \frac{d}{du}\delta\mathbf{x}(u) = {} & (\partial_{\mathbf{x}}f(u))\delta\mathbf{x}(u) \\ & + (\partial_{\boldsymbol{\alpha}}f(u))\delta\boldsymbol{\alpha}(u) \quad \text{with} \quad \delta\mathbf{x}(0) = \delta\mathbf{x}_0(\mathbf{t}(1)) \\ \frac{d}{du}\delta\mathbf{t}(u) = {} & (\partial_{\mathbf{t}}g(u))\delta\mathbf{t}(u) \\ & + (\partial_{\boldsymbol{\beta}}g(u))\delta\boldsymbol{\beta}(u) \quad \text{with} \quad \delta\mathbf{t}(0) = 0 \end{aligned} \tag{62}$$

We introduce the flows P_{su} and R_{su} for all $s, u \in [0,1]$, which are solution of the homogeneous equations:

$$\begin{aligned} \frac{d}{du}P_{su} &= P_{su}(\partial_{\mathbf{x}}f(u)) \quad \text{with} \quad P_{uu} = \mathrm{Id} \\ \frac{d}{du}R_{su} &= R_{su}(\partial_{\mathbf{t}}g(u)) \quad \text{with} \quad R_{uu} = \mathrm{Id} \end{aligned} \tag{63}$$

The method of variations of the parameters leads to the following solution of the ODEs:

$$\begin{aligned} \delta\mathbf{x}(u) &= P_{0u}\delta\mathbf{x}(0) + \int_0^u P_{su}\partial_{\boldsymbol{\alpha}}f(s)\delta\boldsymbol{\alpha}(s)ds \\ \delta\mathbf{t}(u) &= \int_0^u R_{su}\partial_{\boldsymbol{\beta}}g(s)\delta\boldsymbol{\beta}(s)ds \end{aligned} \tag{64}$$

where the variations of the initial condition $\delta\mathbf{x}(0) = \delta\mathbf{x}_0(\mathbf{t}(1))$ equals:

$$\delta\mathbf{x}(0) = (d_{\mathbf{t}(1)}\mathbf{x}_0)\delta\mathbf{t}(1) = (d_{\mathbf{t}(1)}\mathbf{x}_0)\int_0^1 R_{u1}\partial_{\boldsymbol{\beta}}g(u)\delta\boldsymbol{\beta}(u)du, \tag{65}$$

according to (64).

Plugging (64) into (61) leads to the variation of the criterion (noticing that for any L^2 function $F(s,u)$ we have that $\int_0^1\int_0^u F(s,u)dsdu = \int_0^1\int_s^1 F(s,u)duds = \int_0^1\int_u^1 F(u,s)dsdu$):

$$\delta E = \underbrace{\Big((d_{\mathbf{x}(1)}A)P_{01} + \int_0^1 \partial_{\mathbf{x}} L^{\phi}(u) P_{0u} du\Big)}_{\boldsymbol{\eta}(0)^t} \delta\mathbf{x}(0) + \int_0^1 \Big(\partial_{\boldsymbol{\alpha}} L^{\phi}(u) + \underbrace{\Big((d_{\mathbf{x}(1)}A)P_{u1} + \int_u^1 \partial_{\mathbf{x}} L^{\phi}(s) P_{us} ds\Big)}_{\boldsymbol{\eta}(u)^t} \partial_{\boldsymbol{\alpha}} f(u)\Big) \delta\boldsymbol{\alpha}(u) du + \int_0^1 \Big(\partial_{\boldsymbol{\beta}} L^{\psi}(u) + \int_u^1 \partial_{\mathbf{t}} L^{\psi}(s) R_{us} ds\, \partial_{\boldsymbol{\beta}} g(u)\Big) \delta\boldsymbol{\beta}(u) du \tag{66}$$

Now, we denote,

$$\boldsymbol{\eta}(u)^t = (d_{\mathbf{x}(1)}A)P_{u1} + \int_u^1 \partial_{\mathbf{x}} L^{\phi}(s) P_{us} ds \tag{67}$$

which appears twice in (66) as $\boldsymbol{\eta}(0)$ and $\boldsymbol{\eta}(u)$. Given the expression of $\delta\mathbf{x}(0)$ in (65), we have:

$$\delta E = \int_0^1 \big(\partial_{\boldsymbol{\alpha}} L^{\phi}(u) + \eta(u)^t \partial_{\boldsymbol{\alpha}} f(u)\big) \delta\boldsymbol{\alpha}(u) du + \int_0^1 \Big(\partial_{\boldsymbol{\beta}} L^{\psi}(u) + \underbrace{\Big(\boldsymbol{\eta}(0)^t (d_{\mathbf{t}(1)}\mathbf{x}_0) R_{u1} + \int_u^1 \partial_{\mathbf{t}} L^{\psi}(s) R_{us} ds\Big)}_{\boldsymbol{\xi}(u)^t} \partial_{\boldsymbol{\beta}} g(u)\Big) \delta\boldsymbol{\beta}(u) du \tag{68}$$

Denoting

$$\boldsymbol{\xi}(u)^t = \boldsymbol{\eta}(0)^t (d_{\mathbf{t}(1)}\mathbf{x}_0) R_{u1} + \int_u^1 \partial_{\mathbf{t}} L^{\psi}(s) R_{us} ds, \tag{69}$$

we end up with the gradient of the criterion with respect to the L^2 metric written as:

$$\begin{cases} \nabla_{\boldsymbol{\alpha}} E(u) = \partial_{\boldsymbol{\alpha}} L^{\phi}(u)^t + \partial_{\boldsymbol{\alpha}} f(u)^t \boldsymbol{\eta}(u) \\ \nabla_{\boldsymbol{\beta}} E(u) = \partial_{\boldsymbol{\beta}} L^{\psi}(u)^t + \partial_{\boldsymbol{\beta}} g(u)^t \boldsymbol{\xi}(u) \end{cases} \tag{70}$$

The auxiliary variables $\boldsymbol{\eta}(u)$ and $\boldsymbol{\xi}(u)$ depend on the flows R_{us} and P_{us}. Therefore they satisfy a ODE, which we need to make explicit now. The inverse flows are written in integral form as:

$$P_{us} = \mathrm{Id} + \int_u^s P_{rs} \partial_{\mathbf{x}} f(r) dr \quad R_{us} = \mathrm{Id} + \int_u^s R_{rs} \partial_{\mathbf{t}} g(r) dr, \tag{71}$$

so that the auxiliary variable $\boldsymbol{\eta}(u)$ satisfies:

$$\boldsymbol{\eta}(u) = \nabla_{\mathbf{x}(1)} A + \int_u^1 \partial_{\mathbf{x}} L^{\phi}(s)^t ds + \int_u^1 \partial_{\mathbf{x}} f(s)^t (P_{s1})^t \times (\nabla_{\mathbf{x}(1)} A) ds + \int_u^1 \int_u^s (\partial_{\mathbf{x}} f(r))^t (P_{rs})^t (\partial_{\mathbf{x}} L^{\phi}(s))^t dr ds \tag{72}$$

where we denote $\nabla_{\mathbf{x}} A = (d_{\mathbf{x}} A)^t$ for any scalar function A.

Since we have for any L^2 functions $F(r,s)$, $\int_u^1 \int_u^s F(r,s)\, dr ds = \int_u^1 \int_s^1 F(s,r) dr ds$ by permuting the two integrals, we have:

$$\boldsymbol{\eta}(u) = \nabla_{\mathbf{x}(1)} A + \int_u^1 \partial_{\mathbf{x}} L^{\phi}(s)^t + \partial_{\mathbf{x}} f(s)^t \underbrace{\Big((P_{s1})^t \nabla_{\mathbf{x}(1)} A + \int_s^1 (P_{sr})^t (\partial_{\mathbf{x}} L^{\phi}(r))^t dr\Big)}_{\boldsymbol{\eta}(s)} ds \tag{73}$$

The term in the brackets is exactly $\boldsymbol{\eta}(s)$, so that the integral equation satisfied by $\boldsymbol{\eta}(u)$ is eventually given by:

$$\boldsymbol{\eta}(u) = \nabla_{\mathbf{x}(1)} A + \int_u^1 (\partial_{\mathbf{x}} L^{\phi}(s))^t + (\partial_{\mathbf{x}} f(s))^t \boldsymbol{\eta}(s) ds. \tag{74}$$

Similar computations using the integral form of the flow R_{us} leads to the integral equation satisfied by $\boldsymbol{\xi}(u)$:

$$\boldsymbol{\xi}(u) = (d_{\mathbf{t}(1)}\mathbf{x}_0)^t \boldsymbol{\eta}(0) + \int_u^1 (\partial_{\mathbf{t}} L^{\psi}(s))^t + (\partial_{\mathbf{t}} g(s))^t \boldsymbol{\xi}(s) ds. \tag{75}$$

B.3 Gradient in Coordinates

Given the definition of the functions f, g, L^{ϕ} and L^{ψ}, we have:

$$\begin{aligned} &\partial_{\mathbf{x}} f = (\partial_1 + \partial_2)(\mathbf{K}^{\phi}(\mathbf{x},\mathbf{x})\boldsymbol{\alpha}) && \partial_{\boldsymbol{\alpha}} f = \mathbf{K}^{\phi}(\mathbf{x},\mathbf{x}) \\ &\partial_{\mathbf{t}} g = (\partial_1 + \partial_2)(\mathbf{K}^{\psi}(\mathbf{t},\mathbf{t})\boldsymbol{\beta}) && \partial_{\boldsymbol{\beta}} g = \mathbf{K}^{\psi}(\mathbf{t},\mathbf{t}) \\ &\partial_{\mathbf{x}} L^{\phi} = \gamma_{\phi} \boldsymbol{\alpha}^t \Big((\partial_1 + \partial_2)\mathbf{K}^{\phi}(\mathbf{x},\mathbf{x})\alpha\Big) && \partial_{\boldsymbol{\alpha}} L^{\phi} = 2\gamma_{\phi} \boldsymbol{\alpha}^t \mathbf{K}^{\phi}(\mathbf{x},\mathbf{x}) \\ &\partial_{\mathbf{t}} L^{\psi} = \gamma_{\psi} \boldsymbol{\beta}^t \Big((\partial_1 + \partial_2)\mathbf{K}^{\psi}(\mathbf{t},\mathbf{t})\boldsymbol{\beta}\Big) && \partial_{\boldsymbol{\beta}} L^{\psi} = 2\gamma_{\psi} \boldsymbol{\beta}^t \mathbf{K}^{\psi}(\mathbf{t},\mathbf{t}) \end{aligned} \tag{76}$$

so that the gradient with respect to the Sobolev metric (the matrices $\mathbf{K}^{\phi}(\mathbf{x},\mathbf{x})$ and $\mathbf{K}^{\psi}(\mathbf{t},\mathbf{t})$ factorize in (70)) is given as:

$$\begin{cases} \nabla_{\boldsymbol{\alpha}} E(u) = 2\gamma_{\phi}\boldsymbol{\alpha}(u) + \boldsymbol{\eta}(u) \\ \nabla_{\boldsymbol{\beta}} E(u) = 2\gamma_{\psi}\boldsymbol{\beta}(u) + \boldsymbol{\xi}(u) \end{cases} \tag{77}$$

where

$$\boldsymbol{\eta}(u) = \nabla_{\mathbf{x}(1)} A + \int_u^1 \Big((\partial_1+\partial_2)\mathbf{K}^{\phi}(\mathbf{x}(s),\mathbf{x}(s))\alpha(s)\Big)^t (\gamma_{\phi}\boldsymbol{\alpha}(s)+\boldsymbol{\eta}(s))ds \tag{78}$$

and

$$\boldsymbol{\xi}(u) = (d_{\mathbf{t}(1)}\mathbf{x}_0)^t\boldsymbol{\eta}(0) + \int_u^1 \Big((\partial_1+\partial_2)\mathbf{K}^{\psi}(\mathbf{t}(s),\mathbf{t}(s))\boldsymbol{\beta}(s)\Big)^t (\gamma_{\psi}\boldsymbol{\beta}(s)+\boldsymbol{\xi}(s)) \tag{79}$$

The $3NN_{\text{target}}$ vector $\nabla_{\mathbf{x}(1)}A$ is the concatenation of the vectors $\nabla_{x_p(t_j(1))}A$ for $p = 1,\dots,N$ and $j = 1,\dots,N_{\text{target}}$. Similarly, the $3NN_{\text{target}}$ vector is the concatenation of the vectors $\{x_p(t_j(1))\}_{p=1,\dots,N,\,j=1,\dots,N_{\text{target}}}$. $d_{\mathbf{t}(1)}\mathbf{x}_0$ is the $3NN_{\text{target}}$-by-N_{target} matrix: $(d_{t_1(1)}\mathbf{x}_0,\dots,d_{t_{N_{\text{target}}}}\mathbf{x}_0)$. In the vector $d_{t_j(1)}\mathbf{x}_0$ almost every coordinate vanishes except the ones corresponding at the j*th* block of size $3N$: $(d_{t_j(1)}x_1(t_j(1)),\dots,d_{t_j(1)}x_N(t_j(1)))$ (since $d_{t_j(1)}x_p(t_i(1)) = 0$ when $i \neq j$). Therefore, we have: $(d_{t_j(1)}\mathbf{x}_0)^t\eta = \sum_{p=1}^N \left(\frac{dx_p}{dt}(t_j(1))\right)^t \eta_{p,j}$, which is the j*th* coordinate of the N_{target} vector $(d_{\mathbf{t}(1)}\mathbf{x}_0)^t\eta$.

Eventually, using the generic expression (57) for scalar kernels $K^{\phi}(x,y) = k^{\phi}(x,y)\mathrm{Id}$ and $K^{\psi}(x,y) = k^{\psi}(x,y)\mathrm{Id}$, the evolution of $\boldsymbol{\eta}(u)$ and $\boldsymbol{\xi}(u)$ in (78) are written in coordinates as in (28) and (29).

C Algorithms

Algorithm 1 Temporal shape regression

Input:
A set of time-indexed shapes $\{(S_j,t_j)\}$
A baseline $S_0 = \{x_p\}_{p=1,\dots,N}$
A discretization of the interval $[0,T]$: $t_0 = 0,\dots,t_{N_{\text{time}}} = T$

Initialization:
for all $p = 1,\dots,N$, for all $n = 0,\dots,N_{\text{time}}$, $\alpha_p(t_n) \leftarrow 0$

{Gradient descent}
repeat
  {Compute positions of the moving baseline (forward integration)}
  $x_p(t_0) \leftarrow x_p$
  for $n = 0,\dots,N_{\text{time}}-1$ **do**
    for $i = 1,\dots,N$ **do**
      $v = \sum_{q=1}^N K^{\chi}(x_p(t_n),x_q(t_n))\alpha_q(t_n)$
      $x_p(t_{n+1}) \leftarrow x_p(t_n) + v$
    end for
  end for

  {Compute Gradient (backward integration)}
  $\eta_p(t_{N_{\text{time}}}) \leftarrow 0$
  for $n = N_{\text{time}},\dots,1$ **do**
    if t_n is one of the t_j (time-points associated to the shape S_j) **then**
      for $p = 1,\dots,N$ **do**
        $\eta_p(t_n) \leftarrow \eta_p(t_n) + \nabla_p A_n$ (gradient of the matching term)
      end for
    end if
    for $p = 1,\dots,N$ **do**
      $v^{\eta} = \sum_{q=1}^N \Big(\alpha_p(t_n)^t\eta_q(t_n) + \alpha_q(t_n)^t\eta_p(t_n)$
        $+2\gamma^{\chi}\alpha_p(t_n)^t\alpha_q(t_n)\Big)\nabla_1 k^{\chi}(x_p(t_n),x_q(t_n))$
      $\eta_p(t_{n-1}) \leftarrow \eta_p(t_n) + v^{\eta}$
    end for
  end for

  {Update momenta α according to the gradient}
  for $n = 0,\dots,N_{\text{time}}$ **do**
    for $p = 1,\dots,N$ **do**
      $\alpha_p(t_n) \leftarrow \alpha_p(t_n) - \tau\left(2\gamma^{\chi}\alpha_p(t_n) + \eta_p(t_n)\right)$
    end for
  end for
until convergence

Output: the shape evolution $x_p(t_n)$.

Algorithm 2 Spatiotemporal Registration

Input:
A source growth scenario $x_p(t_n)$ for $p = 1, \ldots, N$ and $n = 1, \ldots, N_{\text{time}}$
Target shapes U_j associated to time-points t_j, $j = 1, \ldots, N_{\text{target}}$
Discretization of the interval [0, 1]: $u_0 = 0, \ldots, u_{N_u} = 1$

Initialization:
for $p = 1, \ldots, N$, **for** $j = 1, \ldots, N_{\text{target}}$ **do** $\alpha_{p,j} \leftarrow 0$ **end for end for**
for $j = 1, \ldots, N_{\text{target}}$ **do** $\beta_j \leftarrow 0$ **end for**

{Gradient descent}
repeat
  {Compute spatiotemporal deformation of the source (forward integration: time then space)}
  $t_j(u_0) \leftarrow t_j$
  for $k = 0, \ldots, N_u - 1$ **do**
    for $j = 1, \ldots, N_{\text{target}}$ **do**
      $v = \sum_{i=1}^{N_{\text{target}}} K^{\psi}(t_j(u_k), t_i(u_k))\beta_i(u_k)$
      $t_j(u_{k+1}) \leftarrow t_j(u_k) + v$
    end for
  end for
  $x_{p,j}(u_0) \leftarrow x_p(t_j(u_{N_u}))$
  for $k = 0, \ldots, N_u - 1$ **do**
    for $j = 1, \ldots, N_{\text{target}}, p = 1, \ldots, N$ **do**
      $v = \sum_{i=1}^{N_{\text{target}}} \sum_{q=1}^{N} K^{\phi}(x_{p,j}(u_k), x_{q,i}(u_k))\alpha_{q,i}(u_k)$
      $x_{p,j}(u_{k+1}) \leftarrow x_{p,j}(u_k) + v$
    end for
  end for

  {Compute Gradient (backward integration: space then time)}
  $\eta_{i,p}(u_{N_u}) \leftarrow \nabla_{x_{p,i}(1)} A$ {Gradient of the matching term}
  for $k = N_u, \ldots, 1$ **do**
    for $p = 1, \ldots, N$, $i = 1, \ldots, N_{\text{target}}$ **do**
      $v^{\eta} = \sum_{q=1}^{N} \sum_{j=1}^{N_{\text{target}}} \Big(\alpha_{p,i}(u_k)^t \eta_{q,j}(u_k) + \eta_{p,i}(u_k)^t \alpha_{q,j}(u_k) + 2\gamma^{\phi} \alpha_{p,i}(u_k)^t \alpha_{q,j}(u_k)\Big) \nabla_1 k^{\phi}(x_{p,i}(u_k), x_{q,j}(u_k))$
      $\eta_{i,p}(u_{k-1}) \leftarrow \eta_{i,p}(u_k) + v^{\eta}$
    end for
  end for
  $\xi_j(u_{N_u}) \leftarrow \sum_{p=1}^{N} \frac{dx_p(t)}{dt}(t_j(u_{N_u}))^t \eta_{p,j}(u_0)$
  for $k = N_u, \ldots, 1$ **do**
    for $j = 1, \ldots, N_{\text{target}}$ **do**
      $v^{\xi} \leftarrow \sum_{i=1}^{N_{\text{target}}} \Big(\beta_j(u_k)^t \xi_i(u_k) + \xi_j(u_k)^t \beta_i(u_k) + 2\gamma^{\psi} \beta_j(u_k)^t \beta_i(u_k)\Big) \nabla_1 k^{\psi}(t_j(u_k), t_i(u_k))$
      $\xi_j(u_{k-1}) \leftarrow \xi_j(u_k) + v^{\xi}$
    end for
  end for

  {Update momenta α and β according to the gradient}
  for $k = 0, \ldots, N_u$ **do**
    for $p = 1, \ldots, N$, $j = 1, \ldots, N_{\text{target}}$ **do**
      $\alpha_{p,j}(u_k) \leftarrow \alpha_{p,j}(u_k) - \tau\left(2\gamma^{\phi}\alpha_{p,j}(u_k) + \eta_{p,j}(u_k)\right)$
      $\beta_j(u_k) \leftarrow \beta_j(u_k) - \tau\left(2\gamma^{\psi}\beta_j(u_k) + \xi_j(u_k)\right)$
    end for
  end for
until convergence

Output:
the registered source shapes $\phi(x_p(\psi(t_j))) = x_{p,j}(u_{N_u})$
the parameterization of the morphological deformation $(x_{p,j}(u_k), \alpha_{p,j}(u_k))$
the parameterization of the time warp $(t_j(u_k), \beta_j(u_k))$

Algorithm 3 Spatiotemporal Atlas Construction

Input: A set of time-indexed shapes $\{(S_j^s, t_j^s)\}_{j=1,\ldots,N_{\text{target}}^s}^{s=1,\ldots,N_{\text{subj}}}$, where S_j^s is the j*th* scan (out of N_{target}^s) of the s*th* subject (out of N_{subj}) at age t_j^s.

$M_0 \leftarrow \frac{1}{N_{\text{subj}}} \sum_{s=1}^{N_{\text{subj}}} S_1^s$
$M(t) \leftarrow \chi_t(M_0)$ the regression of every shapes S_j^s at time points t_j^s (using Algorithm 1)
repeat
  {Template-to-subject registration}
  for $s = 1 \ldots N_{\text{subj}}$ **do**
    $(\phi^s, \psi^s) \leftarrow$ spatiotemporal registration of $M(t)$ to S_j^s for $j = 1, \ldots, N_{\text{subj}}^s$ (using Algorithm 1)
  end for

  {Center the template}
  $\Phi_{s,j} \leftarrow \phi^s \circ \chi_{\psi^s(t_j^s)}$ for all $s = 1, \ldots, N_{\text{subj}}$ and $j = 1, \ldots, N_{\text{target}}^s$
  $M_0 \leftarrow$ CenterTemplate$(M_0, \{\Phi_{s,j}\}, \{S_j^s\})$ (Algorithm 4 of Durrleman (2010))

  {Update the mean scenario}
  $M(t) \leftarrow \chi_t(M_0)$ the regression of every shapes $(\phi^s)^{-1}(S_j^s)$ at time points $\psi^s(t_j^s)$ (using Algorithm 1)
until convergence

Output:
One mean scenario $M(t) = \chi_t(M_0)$
N_{subj} spatiotemporal deformations (ϕ^s, ψ^s)

References

Aljabar, P., Bhatia, K., Murgasova, M., Hajnal, J., Boardman, J., Srinivasan, L., et al. (2008). Assessment of brain growth in early childhood using deformation-based morphometry. *NeuroImage*, *39*(1), 348–358.

Allassonnière, S., & Kuhn, E. (2009). Stochastic algorithm for bayesian mixture effect template estimation. *ESAIM Probability and Statistics* (in press).

Chandrashekara, R., Rao, A., Sanchez-Ortiz, G., Mohiaddin, R. H., & Rueckert, D. (2003). Construction of a statistical model for cardiac motion analysis using nonrigid image registration. In *Information processing in medical imaging* (Vol. 2732, pp. 599–610). Springer, Lecture Notes in Computer Science.

Courchesne, E., Campbell, K., & Solso, S. (2011). Brain growth across the life span in autism: Age-specific changes in anatomical pathology. *Brain Research*, *1380*, 138–145 (the Emerging Neuroscience of Autism Spectrum Disorders).

de Craene, M., Camara, O., Bijnens, B. H., & Frangi, A. F. (2009). Large diffeomorphic FFD registration for motion and strain quantification from 3D-US sequences. In *Proceedings of functional imaging and modeling of the heart* (Vol. 5528, pp. 437–446). Spinger, LNCS.

Davis, B., Fletcher, P., Bullitt, E., & Joshi, S. (2007). Population shape regression from random design data. In *Proceedings of international conference on computer vision (ICCV)*, pp. 1–7.

Declerck, J., Feldman, J., & Ayache, N. (1998). Definition of a 4D continuous planispheric transformation for the tracking and the analysis of LV motion. *Medical Image Analysis*, *4*(1), 1–17.

de Waal, F. B. M. (1995). Bonobo sex and society. *Scientific American*, *272*, 82–88.

Dupuis, P., Grenander, U., & Miller, M. (1998). Variational problems on flows of diffeomorphisms for image matching. *Quaterly of Applied Mathematics*, *56*(3), 587–600.

Durrleman, S. (2010). *Statistical models of currents for measuring the variability of anatomical curves, surfaces and their evolution*. Thèse de sciences (phd thesis), Université de Nice-Sophia Antipolis.

Durrleman, S., Pennec, X., Trouvé, A., & Ayache, N. (2009a). Statistical models of sets of curves and surfaces based on currents. *Medical Image Analysis*, *13*(5), 793–808.

Durrleman, S., Pennec, X., Trouvé, A., Gerig, G., & Ayache, N. (2009b) Spatiotemporal atlas estimation for developmental delay detection in longitudinal datasets. In *Medical image computing and computer-assisted intervention—MICCAI* (Vol. 5761, pp. 297–304). Springer, LNCS.

Durrleman, S., Fillard, P., Pennec, X., Trouvé, A., & Ayache, N. (2011). Registration, atlas estimation and variability analysis of white matter fiber bundles modeled as currents. *NeuroImage*, *55*(3), 1073–1090.

Ehrhardt, J., Werner, R., Schmidt-Richberg, A., Schulz, B., & Handels, H. (2008). Generation of a mean motion model of the lung using 4D-CT image data. In *Proceedings of eurographics workshop on visual computing for biomedicine* (pp. 69–76). Eurographics Association.

Fishbaugh, J., Durrleman, S., & Gerig, G. (2011). Estimation of smooth growth trajectories with controlled acceleration from time series shape data. In *Medical image computing and computer-assisted intervention—MICCAI*. Springer, LNCS (to appear).

Gerber, S., Tasdizen, T., Fletcher, T. P., & Whitaker, R. (2010). Manifold modeling for brain population analysis. *Medical Image Analysis 14*(5), 643–653 (special Issue on the 12th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI) 2009).

Gerig, G., Davis, B., Lorenzen, P., Xu, S., Jomier, M., Piven, J., & Joshi, S. (2006). Computational anatomy to assess longitudinal trajectory of brain growth. In *Third international symposium on 3D data processing, visualization, and transmission* (pp. 1041–1047).

Glaunès, J. (2005). *Transport par difféomorphismes de points, de mesures et de courants pour la comparaison de formes et l'anatomie numérique*. PhD thesis, Université Paris 13, http://cis.jhu.edu/joan/TheseGlaunes.pdf.

Gogtay, N., Lu, A., Leow, A., Klunder, A., Lee, A., Chavez, A., et al. (2008). 3D growth pattern abnormalities visualized in childhood-onset schizophrenia using tensor-based morphometry. *Proceedings of the National Academy of Sciences*, *105*(41), 15979–15984.

Grenander, U., Srivastava, A., & Saini, S. (2007). A pattern-theoretic characterization of biological growth. *IEEE Transactions on Medical Imaging*, *26*(5), 648–659.

Hart, G., Shi, Y., Zhu, H., Sanchez, M., Styner, M., & Niethammer, M. (2010). DTI longitudinal atlas construction as an average of growth models. In *Proceedings of international workshop on spatio-temporal image analysis for longitudinal and time-series image data*.

Hazlett, H., Poe, M., Gerig, G., Smith, R., Provenzale, J., Ross, A., et al. (2005). Magnetic resonance imaging and head circumference study of brain size in autism. *The Archives of General Psychiatry*, *62*, 1366–1376.

Hazlett, H., Poe, M., Styner, M., Chappell, C., Smith, R., Vachet, C., et al. (2011). Early brain overgrowth in autism associated with an increase in cortical surface area before age 2 years. *Journal of Archives of General Psychiatry*, *68*(5), 467–476.

Jian, B., & Vemuri, B. C. (2005). A robust algorithm for point set registration using mixture of Gaussians. In: *10th IEEE international conference on computer vision (ICCV 2005)*, 17–20 October 2005, Beijing, China, pp. 1246–1251. http://gmmreg.googlecode.com.

Joshi, S., & Miller, M. (2000). Landmark matching via large deformation diffeomorphisms. *IEEE Transaction on Image Processing*, *9*(8), 1357–1370.

Kano, T. (1992). *The last ape: Pygmy chimpanzee behavior and ecology*. Stanford: Stanford University Press.

Khan, A., & Beg, M. (2008). Representation of time-varying shapes in the large deformation diffeomorphic framework. In *5th IEEE international symposium on biomedical imaging ISBI* (pp. 1521–1524).

Kinzey, W. G. (1984). *The dentition of the pygmy chimpanzee, Pan paniscus*. New York: Plenum.

Kuroda, S. (1989). *Developmental retardation and behavioural characteristics of pygmy chimpanzees*. Cambridge: Harvard University Press.

Mansi, T., Durrleman, S., Bernhardt, B., Sermesant, M., Delingette, H., Voigt, I., et al. (2009). A statistical model of right ventricle in tetralogy of fallot for prediction of remodelling and therapy planning. In *Proceedings of medical image computing and computer assisted intervention (MICCAI)* (Vol. 5761, pp. 214–221). Springer, LNCS.

Miller, I. M., Trouvé, A., & Younes, L. (2002). On the metrics and euler-lagrange equations of computational anatomy. *Annual Review of Biomedical Engineering*, *4*, 375–405.

Miller, M., Trouvé, A., & Younes, L. (2006). Geodesic shooting for computational anatomy. *Journal of Mathematical Imaging and Vision*, *24*(2), 209–228.

Pennec, X., Fillard, P., & Ayache, N. (2006). A Riemannian framework for tensor computing. *International Journal of Computer Vision*, *66*(1), 41–66.

Perperidis, D., Mohiaddin, R. H., & Rueckert, D. (2005). Spatio-temporal free-form registration of cardiac MRI sequences. *Medical Image Analysis*, *9*(5), 441–456.

Peyrat, J. M., Delingette, H., Sermesant, M., Pennec, X., Xu, C., & Ayache, N. (2008). Registration of 4D time-series of cardiac images with multichannel diffeomorphic demons. In *Proceedings of medical image computing and computer assisted intervention (MICCAI)* (Vol. 5242, pp. 972–979). Springer, LNCS.

Qiu, A., Younes, L., Miller, M., & Csernansky, J. (2008). Parallel transport in diffeomorphisms distinguishes the time-dependent pattern of hippocampal surface deformation due to healthy aging and the dementia of the Alzheimer's type. *NeuroImage*, *40*, 68–76.

Qiu, A., Albert, M., Younes, L., & Miller, M. I. (2009). Time sequence diffeomorphic metric mapping and parallel transport track time-dependent shape changes. *NeuroImage*, *45*(1 Supplement 1), S51–S60.

Shea, B. (1989). Heterochrony in human evolution: The case for neoteny reconsidered. *Yearbook of Physical Anthropology*, *32*, 69–101.

Thompson, P. M., Giedd, J. N., Woods, R. P., MacDonald, D., Evans, A. C., & Toga, A. W. (2000). Growth patterns in the developing human brain detected by using continuum-mechanical tensor maps. *Nature*, *404*, 6774.

Trouvé, A. (1998). Diffeomorphisms groups and pattern matching in image analysis. *International Journal of Computer Vision*, *28*(3), 213–221.

Trouvé, A., Vialard, F. X. (2010). Shape splines and stochastic shape evolutions: A second order point of view. *Quaterly of Applied Mathematics* (to appear). http://arxiv.org/abs/1003.3895.

Vaillant, M., Miller, M., Younes, L., & Trouvé, A. (2004). Statistics on diffeomorphisms via tangent space representations. *NeuroImage*, *23*, 161–169.

Xie, Y., Ho, J., & Vemuri, B. C. (2010). Image atlas construction via intrinsic averaging on the manifold of images. In *The twenty-third IEEE conference on computer vision and pattern recognition, CVPR 2010* (pp. 2933–2939). San Francisco, CA, USA: IEEE.

Yushkevich, P. A., Piven, J., Ho, S., Gee, J. C., & Gerig, G. (2006). User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. *Neuroimage*, *31*(3), 1116–1128.